

Compressing rank-structured matrices via randomized sampling

Per-Gunnar Martinsson, Dept. of Applied Math., Univ. of Colorado Boulder

March 22, 2015

Abstract: Randomized sampling has recently been proven a highly efficient technique for computing approximate factorizations of matrices that have low numerical rank. This paper describes an extension of such techniques to a wider class of matrices that are not themselves rank-deficient, but have off-diagonal blocks that are; specifically, the classes of so called *Hierarchically Off-Diagonal Low Rank (HODLR)* matrices and *Hierarchically Block Separable (HBS)* matrices (a.k.a. “Hierarchically Semi-Separable (HSS)” matrices). Such matrices arise frequently in numerical analysis and signal processing, in particular in the construction of fast methods for solving differential and integral equations numerically. These structures admit algebraic operations (matrix-vector multiplications, matrix factorizations, matrix inversion, etc.) to be performed very rapidly; but only once a data-sparse representation of the matrix has been constructed. How to rapidly compute this representation in the first place is much less well understood. The present paper demonstrates that if an $N \times N$ matrix can be applied to a vector in $O(N)$ time, and if the ranks of the off-diagonal blocks are bounded by an integer k , then the cost for constructing a HODLR representation is $O(k^2 N (\log N)^2)$, and the cost for constructing an HBS representation is $O(k^2 N \log N)$ (assuming of course, that the matrix is compressible in the respective format). The point is that when legacy codes (based on, e.g., the Fast Multipole Method) can be used for the fast matrix-vector multiply, the proposed algorithm can be used to obtain the data-sparse representation of the matrix, and then well-established techniques for HODLR/HBS matrices can be used to invert or factor the matrix. The proposed scheme is also useful in simplifying the implementation of certain operations on rank-structured matrices such as the matrix-matrix multiplication, low-rank update, addition, etc.

1. INTRODUCTION

A ubiquitous task in computational science is to rapidly perform linear algebraic operations involving very large matrices. Such operations typically exploit special “structure” in the matrix since the costs of standard techniques tend to scale prohibitively fast with matrix size; for a general $N \times N$ matrix, it costs $O(N^2)$ operations to perform a matrix-vector multiplication, $O(N^3)$ operations to perform Gaussian elimination or to invert the matrix, etc. A well-known form of “structure” in a matrix is sparsity. When at most a few entries in each row of the matrix are non-zero (as is the case, e.g., for matrices arising upon the discretization of differential equations, or representing the link structure of the World Wide Web) matrix-vector multiplications can be performed in $O(N)$ operations instead of $O(N^2)$. The description “data-sparse” applies to a matrix that may be dense, but that shares the key characteristic of a sparse matrix that some linear algebraic operations, typically the matrix-vector multiplication, can to high precision be executed in fewer than $O(N^2)$ operations (often in close to linear time).

Several different formats for rank-structured matrices have been proposed in the literature. In this manuscript, we rely on the so called *Hierarchically Off-Diagonal Low Rank (HODLR)* format. This name was minted in [1], but this class of matrices has a long history. It is a special case of the \mathcal{H} -matrix format introduced by Hackbusch and co-workers [19, 3], and was used explicitly in [30, Sec. 4]. The HODLR format is very easy to describe and easy to use, but can lead to less than optimal performance due to the fact that the basis matrices used to represent large blocks are stored explicitly, leading to a $O(k N \log(N))$ storage requirement for a HODLR matrix whose

off-diagonal blocks have rank at most k . To attain *linear* storage requirements and arithmetic operations, one can switch to a format that expresses all basis matrices *hierarchically*; in other words, the basis matrices used on one level are expressed implicitly in terms of the basis matrices on the next finer level. We sometimes say that we use *nested* basis matrices. To be precise, we use the *Hierarchically Block Separable (HBS)* format that was described in [28, 15]. This format is closely related to the *Hierarchically Semi-Separable (HSS)* [7, 34] format, and is also related to the \mathcal{H}^2 -matrix format [21, 4].

The most straight-forward technique for computing a data-sparse representation of a rank-structured $N \times N$ matrix \mathbf{A} is to explicitly form all matrix elements, and then to compress the off-diagonal blocks using, e.g., the SVD. This approach can be executed stably [35, 20], but it is often prohibitively expensive, with an $O(k N^2)$ asymptotic cost, where k is the rank of the off-diagonal blocks (in the HSS-sense). Fortunately, there exist for specific applications much faster methods for constructing HSS representations. When the matrix \mathbf{A} approximates a boundary integral operator in the plane, the technique of [28] computes a representation in $O(k^2 N)$ time by exploiting representation results from potential theory. In other environments, it is possible to use known regularity properties of the off-diagonal blocks in conjunction with interpolation techniques to obtain rough initial factorizations, and then recompress these to obtain factorizations with close to optimal ranks [4, 29]. A particularly popular version of the “regularity + recompression” method is the so called *Adaptive Cross Approximation* technique which was initially proposed for \mathcal{H} -matrices [2, 5, 23] but has recently been modified to obtain a representation of a matrix in a format similar to the HSS [12].

The purpose of the present paper is to describe a fast and simple randomized technique for computing a data sparse representation of a rank-structured matrix which can rapidly be applied to a vector. The existence of such a technique means that the advantages of the HODLR and HBS formats — fast inversion and factorization algorithms in particular — become available for any matrix that can currently be applied via the FMM, via an \mathcal{H} -matrix calculation, or by any other existing data-sparse format (provided of course that the matrix is in principle rank-structured). In order to describe the cost of the algorithm precisely, we introduce some notation: We let \mathbf{A} be an $N \times N$ matrix whose off-diagonal blocks have maximal rank k , we let T_{mult} denote the time required to perform a matrix-vector multiplication $\mathbf{x} \mapsto \mathbf{A}\mathbf{x}$ or $\mathbf{x} \mapsto \mathbf{A}^*\mathbf{x}$, we let T_{rand} denote the cost of constructing a pseudo random number from a normalized Gaussian distribution, and T_{flop} denote the cost of a floating point operation. The computational cost T_{total} of the algorithm for the HBS format then satisfies

$$T_{\text{total}} \sim T_{\text{mult}} \times k \log(N) + T_{\text{rand}} \times (k + p) N \log(N) + T_{\text{flop}} \times k^2 N \log(N), \quad (1)$$

where p is a tuning parameter that balances computational cost against the probability of not meeting the requested accuracy. Setting $p = 10$ is often a good choice which leads to a “failure probability” of less than 10^{-9} , see Remark 1. In particular, if T_{mult} is $O(N)$, then the method presented here has overall complexity $O(k^2 N \log(N))$. For the HODLR format, an additional factor of $\log N$ arises, cf. (13).

The work presented is directly inspired by [27] (which is based on a 2008 preprint [26]), which presented a similar algorithm with $O(k^2 N)$ complexity for the compression of an HBS matrix. This is better by a factor of $\log(N)$ compared to the present work, but the algorithm of [27] has a serious limitation in that it requires the ability to evaluate $O(k N)$ entries of the matrix to be compressed. This algorithm was later refined by Xia [33] and applied to the task of accelerated the “nested dissection” direct solver [13, 11] for elliptic PDEs to $O(N)$ complexity. In 2011, by L. Lin, J. Lu, and L. Ying presented an alternative algorithm [25] that interacts with the matrix *only* via the matrix-vector multiplication, which makes the randomized compression idea much more broadly applicable than the algorithm in [27]. However, this came at the cost of requiring $O(k \log(N))$ matrix-vector multiplies (just like the present work). The algorithm proposed here is an evolution of [25] that does away with a step of least-square fitting that led to a magnification of the sampling error resulting from the randomized approximation. Moreover, we here present a new strategy for enforcing the “nestedness” of the basis matrices that is required in the HBS format.

Remark 1. *The technique described in this paper utilizes a method for computing approximate low-rank factorizations of matrices that is based on randomized sampling [31, 32, 22]. As a consequence, there is in principle*

a non-zero risk that the method may fail to produce full accuracy in any given realization of the algorithm. This risk can be controlled by the user via the choice of the tuning parameter p in (1), for details see Section 2.5. Moreover, unlike some better known randomized algorithms such as Monte Carlo, the accuracy of the output of the algorithms under discussion here is typically very high; in the environment described in the present paper, approximation errors of less than 10^{-10} are entirely typical.

2. PRELIMINARIES

2.1. Notation. Throughout the paper, we measure vectors in \mathbb{R}^n using their Euclidean norm. The default norm for matrices will be the corresponding operator norm $\|\mathbf{A}\| = \sup_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|$, although we will sometimes also use the Frobenius norm $\|\mathbf{A}\|_{\text{Fro}} = \left(\sum_{i,j} |\mathbf{A}(i,j)|^2\right)^{1/2}$.

We use the notation of Golub and Van Loan [16] to specify submatrices. In other words, if \mathbf{B} is an $m \times n$ matrix with entries b_{ij} , and $I = [i_1, i_2, \dots, i_k]$ and $J = [j_1, j_2, \dots, j_\ell]$ are two index vectors, then we let $\mathbf{B}(I, J)$ denote the $k \times \ell$ matrix

$$\mathbf{B}(I, J) = \begin{bmatrix} b_{i_1 j_1} & b_{i_1 j_2} & \cdots & b_{i_1 j_\ell} \\ b_{i_2 j_1} & b_{i_2 j_2} & \cdots & b_{i_2 j_\ell} \\ \vdots & \vdots & & \vdots \\ b_{i_k j_1} & b_{i_k j_2} & \cdots & b_{i_k j_\ell} \end{bmatrix}.$$

We let $\mathbf{B}(I, :)$ denote the matrix $\mathbf{B}(I, [1, 2, \dots, n])$, and define $\mathbf{B}(:, J)$ analogously.

The transpose of \mathbf{B} is denoted \mathbf{B}^* , and we say that a matrix \mathbf{U} is *orthonormal* if its columns form an orthonormal set, so that $\mathbf{U}^* \mathbf{U} = I$.

2.2. The QR factorization. Any $m \times n$ matrix \mathbf{A} admits a *QR factorization* of the form

$$\begin{array}{ccccc} \mathbf{A} & & \mathbf{P} & = & \mathbf{Q} & \mathbf{R}, \\ & m \times n & n \times n & & m \times r & r \times n \end{array} \quad (2)$$

where $r = \min(m, n)$, \mathbf{Q} is orthonormal, \mathbf{R} is upper triangular, and \mathbf{P} is a permutation matrix. The permutation matrix \mathbf{P} can more efficiently be represented via a vector $J_c \in \mathbb{Z}_+^n$ of column indices such that $\mathbf{P} = \mathbf{I}(:, J_c)$ where \mathbf{I} is the $n \times n$ identity matrix. As a result, the factorization can be written as:

$$\begin{array}{ccccc} \mathbf{A}(:, J_c) & = & \mathbf{Q} & & \mathbf{R}, \\ & m \times n & & m \times r & r \times n \end{array}$$

The QR-factorization is often built incrementally via a greedy algorithm such as column pivoted Gram-Schmidt. This allows one to stop after the first k terms have been computed to obtain a “partial QR-factorization of \mathbf{A} ”:

$$\begin{array}{ccccc} \mathbf{A}(:, J_c) & = & [\mathbf{Q}^{(1)} \quad \mathbf{Q}^{(2)}] & \begin{bmatrix} \mathbf{R}^{(1)} \\ \mathbf{R}^{(2)} \end{bmatrix}, \\ & m \times n & m \times r & r \times n \end{array}$$

That is, taking the first k columns of \mathbf{Q} and the first k rows of \mathbf{R} , we can obtain the approximation:

$$\mathbf{A}(:, J_c) \approx \mathbf{Q}_k \mathbf{R}_k \quad (3)$$

We note that the partial factors \mathbf{Q}_k and \mathbf{R}_k can be obtained after k steps of the pivoted QR algorithm, without having to compute the full matrices \mathbf{Q} and \mathbf{R} .

2.3. The singular value decomposition (SVD). Let \mathbf{A} denote an $m \times n$ matrix, and set $r = \min(m, n)$. Then \mathbf{A} admits a factorization

$$\begin{array}{ccccccccc} \mathbf{A} & = & \mathbf{U} & \Sigma & \mathbf{V}^*, \\ m \times n & & m \times r & r \times r & r \times n \end{array} \quad (4)$$

where the matrices \mathbf{U} and \mathbf{V} are orthonormal, and Σ is diagonal. We let $\{\mathbf{u}_i\}_{i=1}^r$ and $\{\mathbf{v}_i\}_{i=1}^r$ denote the columns of \mathbf{U} and \mathbf{V} , respectively. These vectors are the left and right singular vectors of \mathbf{A} . The diagonal elements $\{\sigma_j\}_{j=1}^r$ of Σ are the singular values of \mathbf{A} . We order these so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$. We let \mathbf{A}_k denote the truncation of the SVD to its first k terms, so that $\mathbf{A}_k = \sum_{j=1}^k \sigma_j \mathbf{u}_j \mathbf{v}_j^*$. It is easily verified that

$$\|\mathbf{A} - \mathbf{A}_k\|_{\text{spectral}} = \sigma_{k+1}, \quad \text{and that} \quad \|\mathbf{A} - \mathbf{A}_k\|_{\text{Fro}} = \left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2 \right)^{1/2},$$

where $\|\cdot\|_{\text{spectral}}$ denotes the operator norm and $\|\cdot\|_{\text{Fro}}$ denotes the Frobenius norm. Moreover, the Eckart-Young theorem states that these errors are the smallest possible errors that can be incurred when approximating \mathbf{A} by a matrix of rank k .

2.4. The interpolatory decomposition (ID). Let \mathbf{A} denote an $m \times n$ matrix of rank k . Then \mathbf{A} admits the factorization

$$\begin{array}{ccccccccc} \mathbf{A} & = & \mathbf{A}(:, J) & \mathbf{X}, \\ m \times n & & m \times k & k \times n \end{array}$$

where J is a vector of indices marking k of the columns of \mathbf{A} , and the $k \times n$ matrix \mathbf{X} has the $k \times k$ identity matrix as a submatrix and has the property that all its entries are bounded by 1 in magnitude. In other words, the interpolative decomposition picks k columns of \mathbf{A} as a basis for the column space of \mathbf{A} and expresses the remaining columns in terms of the chosen ones. The ID can be viewed as a modification to so the called *Rank-Revealing QR factorization* [6]. It can be computed in a stable and accurate manner using the techniques of [18], as described in [8]. (Practical algorithms for computing the interpolative decomposition produce a matrix \mathbf{X} whose elements slightly exceed 1 in magnitude.)

2.5. Randomized compression. Let \mathbf{A} be a given $m \times n$ matrix that can accurately be approximated by a matrix of rank k , and suppose that we seek to determine a matrix \mathbf{Q} with orthonormal columns (as few as possible) such that

$$\|\mathbf{A} - \mathbf{Q} \mathbf{Q}^* \mathbf{A}\|$$

is small. In other words, we seek a matrix \mathbf{Q} whose columns form an approximate orthonormal basis (ON-basis) for the column space of \mathbf{A} . This task can efficiently be solved via the following randomized procedure:

- (1) Pick a small integer p representing how much “over-sampling” we do. ($p = 10$ is often good.)
- (2) Form an $n \times (k + p)$ matrix Ω whose entries are iid normalized Gaussian random numbers.
- (3) Form the product $\mathbf{Y} = \mathbf{A} \Omega$.
- (4) Construct a matrix \mathbf{Q} whose columns form an ON-basis for the columns of \mathbf{Y} .

Note that each column of the “sample” matrix \mathbf{Y} is a random linear combination of the columns of \mathbf{A} . We would therefore expect the algorithm described to have a high probability of producing an accurate result when p is a large number. It is perhaps less obvious that this probability depends only on p (not on m or n , or any other properties of \mathbf{A}), and that it approaches 1 extremely rapidly as p increases. In fact, one can show that the basis \mathbf{Q} determined by the scheme above satisfies

$$\|\mathbf{A} - \mathbf{Q} \mathbf{Q}^* \mathbf{A}\| \leq \left[1 + 11\sqrt{k+p} \cdot \sqrt{\min\{m, n\}} \right] \sigma_{k+1}, \quad (5)$$

with probability at least $1 - 6 \cdot p^{-p}$, see [22, Sec. 1.5]. The error bound (5) indicates that the error produced by the randomized sampling procedure can be larger than the theoretically minimal error σ_{k+1} by a factor of

$1 + 11\sqrt{k + p} \cdot \sqrt{\min\{m, n\}}$. This crude bound is typically very pessimistic; for specific situations sharper bounds have been proved, see [22].

Definition 1. Let \mathbf{A} be an $m \times n$ matrix, and let ε be a positive number. We say that an $m \times \ell$ matrix \mathbf{Y} an ε -spanning matrix for \mathbf{A} , if $\|\mathbf{A} - \mathbf{Y}\mathbf{Y}^\dagger\mathbf{A}\| \leq \varepsilon$. Informally, this means that the columns of \mathbf{Y} span the column space \mathbf{A} to within precision ε . Furthermore, we say that an $m \times \ell$ matrix \mathbf{Q} is an ε -basis matrix for \mathbf{A} if its columns are orthonormal, and if $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \varepsilon$.

2.6. Functions for low-rank factorizations. For future reference, we introduce two functions “qr” and “svd” that can operate in three difference modes. In the first mode, they produce the full (“economy size”) factorizations described in section 2.2 and 2.3, respectively,

$$[\mathbf{Q}, \mathbf{R}, J] = \text{qr}(\mathbf{A}), \quad [\mathbf{U}, \mathbf{D}, \mathbf{V}] = \text{svd}(\mathbf{A}), \quad \text{and} \quad [\mathbf{X}, J] = \text{id}(\mathbf{A}).$$

In practice, we execute these factorizations using standard LAPACK library functions. In the second mode, we provide an integer k and obtain partial factorizations of rank k ,

$$[\mathbf{Q}, \mathbf{R}, J] = \text{qr}(\mathbf{A}, k), \quad [\mathbf{U}, \mathbf{D}, \mathbf{V}] = \text{svd}(\mathbf{A}, k), \quad \text{and} \quad [\mathbf{X}, J] = \text{id}(\mathbf{A}, k).$$

Then the matrices \mathbf{Q} , \mathbf{U} , \mathbf{D} , \mathbf{V} have precisely k columns, and \mathbf{R} has precisely k rows. In the third mode, we provide a real number $\varepsilon \in (0, 1)$ and obtain partial factorizations

$$[\mathbf{Q}, \mathbf{R}, J] = \text{qr}(\mathbf{A}, \varepsilon), \quad [\mathbf{U}, \mathbf{D}, \mathbf{V}] = \text{svd}(\mathbf{A}, \varepsilon), \quad \text{and} \quad [\mathbf{X}, J] = \text{id}(\mathbf{A}, \varepsilon),$$

such that

$$\|\mathbf{A}(:, J) - \mathbf{QR}\| \leq \varepsilon \quad \|\mathbf{A} - \mathbf{UDV}^*\| \leq \varepsilon, \quad \text{and} \quad \|\mathbf{A} - \mathbf{A}(:, J)\mathbf{X}\| \leq \varepsilon.$$

In practice, for a *small* input matrix \mathbf{A} , we execute mode 2 and mode 3 by calling the LAPACK routine for a *full* factorization (the ID can be obtained from the full QR), and then simply truncating the result. If \mathbf{A} is large, then we use the randomized sampling technique of Section 2.5.

Remark 2. The differentiation between modes 2 and 3 for qr and svd is communicated by whether the second argument is an integer (mode 2) or a real number $\varepsilon \in (0, 1)$ (mode 3). This is slightly questionable notation, but it keeps the formulas clean, and hopefully does not cause confusion.

2.7. A binary tree structure. Both the HODLR and the HBS representations of an $M \times M$ matrix \mathbf{A} are based on a partition of the index vector $I = [1, 2, \dots, M]$ into a binary tree structure. We let I form the root of the tree, and give it the index 1, $I_1 = I$. We next split the root into two roughly equi-sized vectors I_2 and I_3 so that $I_1 = I_2 \cup I_3$. The full tree is then formed by continuing to subdivide any interval that holds more than some preset fixed number m of indices. We use the integers $\ell = 0, 1, \dots, L$ to label the different levels, with 0 denoting the coarsest level. A *leaf* is a node corresponding to a vector that never got split. For a non-leaf node τ , its *children* are the two boxes σ_1 and σ_2 such that $I_\tau = I_{\sigma_1} \cup I_{\sigma_2}$, and τ is then the *parent* of σ_1 and σ_2 . Two boxes with the same parent are called *siblings*. These definitions are illustrated in Figure 1. For any node τ , let n_τ denote the number of indices in I_τ .

2.8. The HODLR data sparse matrix format. The *Hierarchically Off-Diagonal Low Rank (HODLR)* property is, as the name implies, a condition that the off-diagonal blocks of a matrix \mathbf{A} should have low (numerical) rank. To be precise, given a hierarchical partitioning of the index vector, cf. Section 2.7, a computational tolerance ε , and a bound on the rank k , we require that for any sibling pair $\{\alpha, \beta\}$, the corresponding block

$$\mathbf{A}_{\alpha, \beta} = \mathbf{A}(I_\alpha, I_\beta)$$

should have ε -rank at most k . The tessellation resulting from the tree in Figure 1 is shown in Figure 2. We then represent each off-diagonal block via a rank- k factorization

$$\begin{array}{rcl} \mathbf{A}_{\alpha, \beta} & = & \mathbf{U}_\alpha \quad \tilde{\mathbf{A}}_{\alpha, \beta} \quad \mathbf{U}_\beta, \\ n_\alpha \times n_\beta & & n_\alpha \times k \quad k \times k \quad k \times n_\beta \end{array}$$

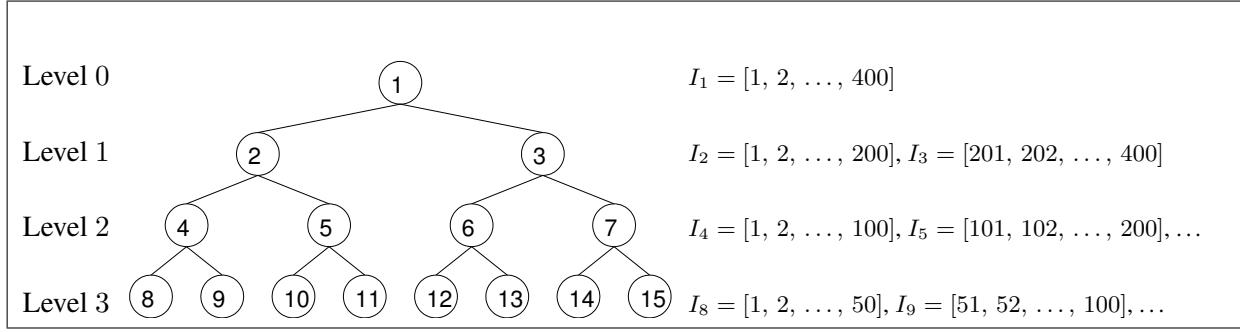


FIGURE 1. Numbering of nodes in a fully populated binary tree with $L = 3$ levels. The root is the original index vector $I = I_1 = [1, 2, \dots, 400]$.

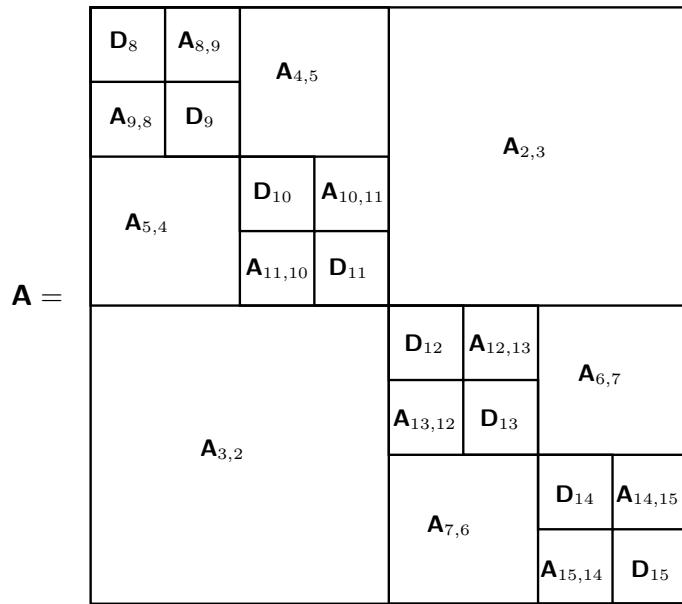


FIGURE 2. A HODLR matrix \mathbf{A} tessellated in accordance with the tree in Figure 1. Every off-diagonal block $\mathbf{A}_{\alpha,\beta}$ that is marked in the figure should have ε -rank at most k .

where \mathcal{U}_α and \mathcal{B}_β are orthonormal matrices. It is easily verified that if we required each leaf node to have at most $O(k)$ nodes, then it takes $O(k N \log N)$ storage to store all factors required to represent \mathbf{A} , and a matrix-vector multiplication can be executed using $O(k N \log N)$ flops.

For future reference, we define for a given HODLR matrix \mathbf{A} a “level-truncated” matrix $\mathbf{A}^{(\ell)}$ as the matrix obtained by zeroing out any block associated with levels finer than ℓ . In other words,

$$\mathbf{A}^{(1)} = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{2,3} \\ \mathbf{A}_{3,2} & \mathbf{0} \end{bmatrix}, \quad \mathbf{A}^{(2)} = \begin{bmatrix} \mathbf{0} & \mathbf{A}_{4,5} & \mathbf{A}_{2,3} \\ \mathbf{A}_{5,4} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{6,7} & \mathbf{0} \\ \mathbf{A}_{3,2} & \mathbf{0} & \mathbf{A}_{7,6} \\ \mathbf{A}_{7,6} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \text{etc.}$$

Remark 3. In our description of rank-structured matrices in sections 2.8 and 2.9, we generally assume that the numerical rank is the same number k for every off-diagonal block. In practice, we typically estimate the ε -rank for any specific off-diagonal block adaptively to save both storage and flop counts.

2.9. The HBS data sparse matrix format. The HODLR format is simple to describe and to use, but is slightly inefficient in that it requires the user to store for every node τ , the basis matrices \mathcal{U}_τ and \mathcal{V}_τ , which can be quite long. The *Hierarchically Block Separable (HBS)* format is designed to overcome this problem by expressing these matrices *hierarchically*. To be precise, suppose that τ is a node which children $\{\alpha, \beta\}$, and that we can find a *short* matrix \mathbf{U}_τ such that

$$\mathcal{U}_\tau = \begin{bmatrix} \mathcal{U}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathcal{U}_\beta \end{bmatrix} \mathbf{U}_\tau. \quad (6)$$

$n_\tau \times k$ $n_\tau \times 2k$ $2k \times k$

The point is that if we have the long basis matrices \mathcal{U}_α and \mathcal{U}_β available for the children, then all we need to store in order to be able to apply \mathcal{U}_τ is the short matrix \mathbf{U}_τ . This process can now be continued recursively. For instance, if $\{\gamma, \delta\}$ are the children of α , and $\{\nu, \mu\}$ are the children of β , we assume there exist matrices \mathbf{U}_α and \mathbf{U}_β such that

$$\mathcal{U}_\alpha = \begin{bmatrix} \mathcal{U}_\gamma & \mathbf{0} \\ \mathbf{0} & \mathcal{U}_\delta \end{bmatrix} \mathbf{U}_\alpha, \quad \text{and} \quad \mathcal{U}_\beta = \begin{bmatrix} \mathcal{U}_\mu & \mathbf{0} \\ \mathbf{0} & \mathcal{U}_\nu \end{bmatrix} \mathbf{U}_\beta.$$

$n_\alpha \times k$ $n_\alpha \times 2k$ $2k \times k$ $n_\beta \times k$ $n_\beta \times 2k$ $2k \times k$

Then \mathcal{U}_τ can be expressed as

$$\mathcal{U}_\tau = \begin{bmatrix} \mathcal{U}_\gamma & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathcal{U}_\delta & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{U}_\mu & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathcal{U}_\nu \end{bmatrix} \begin{bmatrix} \mathbf{U}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_\beta \end{bmatrix} \mathbf{U}_\tau.$$

$n_\tau \times k$ $n_\tau \times 4k$ $4k \times 2k$ $2k \times k$

By continuing this process down to the leaves, it becomes clear that we only need to store the “long” basis matrices for a leaf node (and they are not in fact long for a leaf node!); for every other node, it is sufficient to store the small matrix \mathbf{U}_τ . The process for storing the long basis matrices \mathcal{V}_τ via small matrices \mathbf{V}_τ of size $2k \times k$ is of course exactly analogous.

For a general HODLR matrix, there is no guarantee that a relationship such as (6) need hold. We need to impose an additional condition on the long basis matrices \mathcal{U}_τ and \mathcal{V}_τ . To this end, given a node τ , let us define a *neutered row block* as the off-diagonal block $\mathbf{A}(I_\tau, I_\tau^c)$, where I_τ^c is the complement of I_τ within the vector $[1, 2, 3, \dots, N]$, cf. Figure 3. We then require that the columns of the long basis matrix \mathcal{U}_τ must span the columns of $\mathbf{A}(I_\tau, I_\tau^c)$. Observe that for a node τ with sibling σ , the sibling matrix $\mathbf{A}(I_\tau, I_\sigma)$ is a submatrix of the neutered row block $\mathbf{A}(I_\tau, I_\tau^c)$ since $I_\sigma \subseteq I_\tau^c$. This means that the new requirement of the basis matrices is more restrictive, and that typically the ranks required will be larger for any given precision. However, once the long basis matrices satisfy the more restrictive requirement, it is *necessarily* the case that (6) holds for some small matrix \mathbf{U}_τ .

We analogously define the *neutered column block* for τ as the matrix $\mathbf{A}(I_\tau^c, I_\tau)$ and require that the columns of \mathcal{V}_τ span the rows of $\mathbf{A}(I_\tau^c, I_\tau)$.

Definition 2. We say that a HODLR matrix \mathbf{A} is an HBS matrix if, for every parent node τ with children $\{\alpha, \beta\}$, there exist “small” basis matrices \mathbf{U}_τ and \mathbf{V}_τ such that

$$\mathcal{U}_\tau = \begin{bmatrix} \mathcal{U}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathcal{U}_\beta \end{bmatrix} \mathbf{U}_\tau, \quad \text{and} \quad \mathcal{V}_\tau = \begin{bmatrix} \mathcal{V}_\alpha & \mathbf{0} \\ \mathbf{0} & \mathcal{V}_\beta \end{bmatrix} \mathbf{V}_\tau.$$

$n_\tau \times k$ $n_\tau \times 2k$ $2k \times k$ $n_\tau \times k$ $n_\tau \times 2k$ $2k \times k$

While standard practice is to require all basis matrices \mathcal{U}_τ and \mathcal{V}_τ to be *orthonormal*, we have found that it is highly convenient to use the interpolatory decomposition (ID) to represent the off-diagonal blocks. The key advantage is that then the sibling interaction matrices which be *submatrices* of the original matrix. This improves interpretability, and also slightly reduces storage requirements.

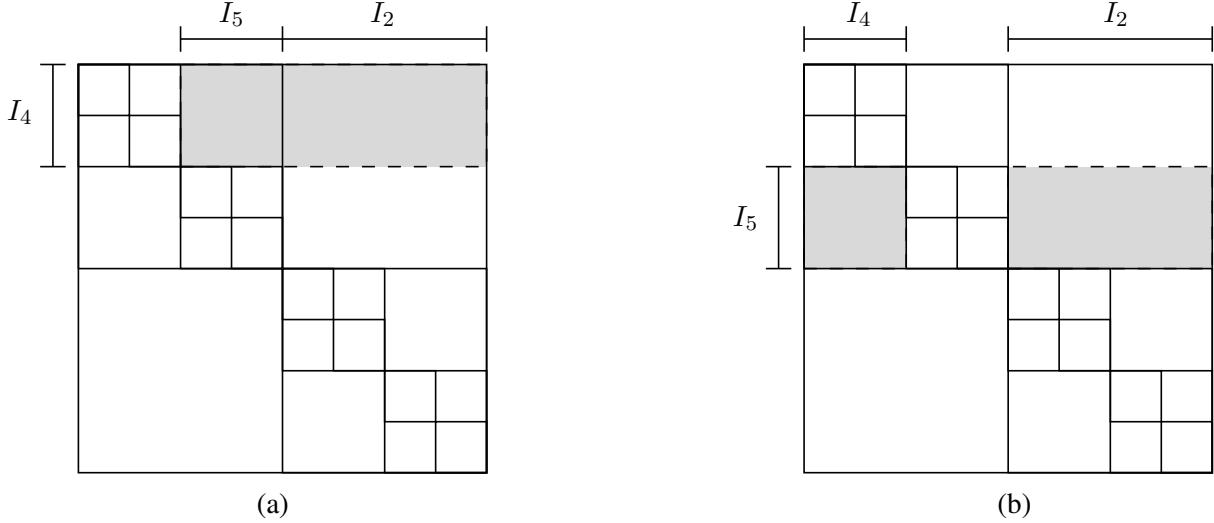


FIGURE 3. Illustration of the *neutered row blocks* for the nodes 4 and 5, with parent 2. (a) The block $\mathbf{A}(I_4, I_4^c)$ is marked in grey. Observe that $\mathbf{A}(I_4, I_4^c) = [\mathbf{A}(I_4, I_5), \mathbf{A}(I_4, I_2)]$. (b) The block $\mathbf{A}(I_5, I_5^c)$ is marked in grey. Observe that $\mathbf{A}(I_5, I_5^c) = [\mathbf{A}(I_5, I_4), \mathbf{A}(I_5, I_2)]$.

Definition 3. We say that a HBS matrix \mathbf{A} is in HBS-ID format if every basis matrix \mathbf{U}_τ and \mathbf{V}_τ contains a $k \times k$ identity matrix, and every sibling interaction matrix $\tilde{\mathbf{A}}_{\alpha, \beta}$ is a submatrix of \mathbf{A} . In other words, there exist some index sets \tilde{I}_α and \hat{I}_β such that

$$\tilde{\mathbf{A}}_{\alpha, \beta} = \mathbf{A}(\tilde{I}_\alpha, \hat{I}_\beta).$$

The index sets \tilde{I}_τ and \hat{I}_τ are called the row skeleton and column skeleton of box τ , respectively. We enforce that these are “nested,” which is to say that if the children of τ are $\{\alpha, \beta\}$, then

$$\tilde{I}_\tau \subseteq \tilde{I}_\alpha \cup \tilde{I}_\beta \quad \text{and} \quad \hat{I}_\tau \subseteq \hat{I}_\alpha \cup \hat{I}_\beta.$$

Remark 4. The straight-forward way to build a basis matrix \mathbf{U}_τ for a node τ is to explicitly form the corresponding neutered row block $\mathbf{A}(I_\tau, I_\tau^c)$ and then compress it (perform column pivoted Gram-Schmidt on its columns to form an ON basis \mathbf{U}_τ , or perform column pivoted Gram-Schmidt on its rows to form the ID). However, suppose that we can somehow construct a smaller matrix \mathbf{Y}_τ of size $n_\tau \times \ell$ with the property that the columns of \mathbf{Y}_τ span the columns of $\mathbf{A}(I_\tau, I_\tau^c)$. Then it would be sufficient to process the columns of \mathbf{Y}_τ to build a basis for $\mathbf{A}(I_\tau, I_\tau^c)$. For instance, if we orthonormalize the columns of \mathbf{Y}_τ to form a basis matrix $\mathbf{U}_\tau = \mathbf{qr}(\mathbf{Y}_\tau)$, then the columns of \mathbf{U}_τ will necessarily form an orthonormal basis for the columns of $\mathbf{A}(I_\tau, I_\tau^c)$. The key point here is that one can often find such a matrix \mathbf{Y}_τ with a small number ℓ of columns. In [28] we use a representation theorem from potential theory to find such a matrix \mathbf{Y}_τ when \mathbf{A} comes from the discretization of a boundary integral equation of mathematical physics. In [27] we do this via randomized sampling, so that $\mathbf{Y}_\tau = \mathbf{A}(I_\tau, I_\tau^c) \mathbf{\Omega}$ for some Gaussian random matrix $\mathbf{\Omega}$. The main point of [27] is that this can be done by applying all of \mathbf{A} to a single random matrix with $N \times \ell$ columns, where $\ell \approx k$. In the current manuscript, we use a similar strategy, but we now require the application of \mathbf{A} to a set of $O(k)$ random vector for each level.

3. AN ALGORITHM FOR COMPRESSING A HODLR MATRIX

The algorithm consists of a single sweep through the levels in the tree of nodes, starting from the root (the entire domain), and processing each level of successively smaller blocks at a time.

In presenting the algorithm, we let \mathbf{G}_τ denote a random matrix of size $n_\tau \times r$ drawn from a Gaussian distribution. Blocked matrices are drawn with the blocks to be processed set in red type. To minimize clutter, blocks of the matrix that will not play a part in the current step are marked with a star (“*”) and may or may not be zero.

Processing level 0 (the root of the tree): Let $\{\alpha, \beta\}$ denote the children of the root node (in our standard ordering, $\alpha = 2$ and $\beta = 3$). Our objective is now to find low-rank factorizations of the off-diagonal blocks

$$\mathbf{A} = \begin{bmatrix} * & \mathbf{A}_{\alpha\beta} \\ \mathbf{A}_{\alpha\beta} & * \end{bmatrix}.$$

To this end, we build two random matrix, each of size $N \times r$, and defined by

$$\Omega_1 = \begin{bmatrix} \mathbf{G}_\alpha \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \Omega_2 = \begin{bmatrix} \mathbf{0} \\ \mathbf{G}_\beta \end{bmatrix}.$$

Then construct the matrices of samples

$$\mathbf{Y}_1 = \mathbf{A}\Omega_2 = \begin{bmatrix} \mathbf{A}_{\alpha\beta}\mathbf{G}_\beta \\ * \end{bmatrix} \quad \text{and} \quad \mathbf{Y}_2 = \mathbf{A}\Omega_1 = \begin{bmatrix} * \\ \mathbf{A}_{\beta\alpha}\mathbf{G}_\alpha \end{bmatrix}.$$

Supported by the results on randomized sampling described in Section 2.5, we now know that it is almost certain that the columns in the top block of \mathbf{Y}_1 will span the columns of $\mathbf{A}_{\alpha\beta}$. By orthonormalizing the columns of $\mathbf{Y}_1(I_\alpha, :)$, we therefore obtain an ON-basis for the column space of $\mathbf{A}_{\alpha\beta}$. In other words, we set

$$\mathcal{U}_\alpha = \text{qr}(\mathbf{Y}_1(I_\alpha, :)), \quad \text{and} \quad \mathcal{U}_\beta = \text{qr}(\mathbf{Y}_2(I_\beta, :)),$$

and then we know that \mathcal{U}_α and \mathcal{U}_β will serve as the relevant basis matrices in the HODLR representation of \mathbf{A} . To compute \mathbf{V}_α , \mathbf{V}_β , $\mathbf{B}_{\alpha\beta}$ and $\mathbf{B}_{\beta\alpha}$ we now need to form the matrices $\mathcal{U}_\alpha^* \mathbf{A}_{\alpha\beta}$ and $\mathcal{U}_\beta^* \mathbf{A}_{\beta\alpha}$. To do this through our “black-box” matrix-matrix multiplier, we form the new test matrices

$$\Omega_1 = \begin{bmatrix} \mathcal{U}_\alpha \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \Omega_2 = \begin{bmatrix} \mathbf{0} \\ \mathcal{U}_\beta \end{bmatrix}.$$

Then construct the matrices of samples

$$\mathbf{Z}_1 = \mathbf{A}^* \Omega_2 = \begin{bmatrix} \mathbf{A}_{\beta\alpha}^* \mathcal{U}_\beta \\ * \end{bmatrix} \quad \text{and} \quad \mathbf{Z}_2 = \mathbf{A}^* \Omega_1 = \begin{bmatrix} * \\ \mathbf{A}_{\alpha\beta}^* \mathcal{U}_\alpha \end{bmatrix}.$$

All that remains is now to compute QR factorizations

$$[\mathbf{V}_\alpha, \mathbf{B}_{\beta\alpha}] = \text{qr}(\mathbf{Z}_1(I_\alpha, :)), \quad \text{and} \quad [\mathbf{V}_\beta, \mathbf{B}_{\alpha\beta}] = \text{qr}(\mathbf{Z}_2(I_\beta, :)). \quad (7)$$

Remark 5. At a slight increase in cost, one can obtain diagonal sibling interaction matrices $\mathbf{B}_{\alpha\beta}$ and $\mathbf{B}_{\beta\alpha}$. We would then replace the QR factorization in (7) by a full SVD

$$[\mathbf{V}_\alpha, \mathbf{B}_{\beta\alpha}, \hat{\mathbf{U}}_\beta] = \text{svd}(\mathbf{Z}_1(I_\alpha, :)), \quad \text{and} \quad [\mathbf{V}_\beta, \mathbf{B}_{\alpha\beta}, \hat{\mathbf{U}}_\alpha] = \text{svd}(\mathbf{Z}_2(I_\beta, :)). \quad (8)$$

This step then requires an update to the long basis matrices for the column space:

$$\mathcal{U}_\alpha \leftarrow \mathcal{U}_\alpha \hat{\mathbf{U}}_\alpha \quad \text{and} \quad \mathcal{U}_\beta \leftarrow \mathcal{U}_\beta \hat{\mathbf{U}}_\beta. \quad (9)$$

Processing level 1: Now that all off-diagonal blocks on level 1 have been computed, we can use this information to compress the blocks on level 2. We let $\{\alpha, \beta, \gamma, \delta\}$ denote the boxes on level 2 (in standard ordering, $\alpha = 4$, $\beta = 5$, $\gamma = 6$, $\delta = 7$). Our objective is now to construct low-rank approximations to the blocks marked in red:

$$\mathbf{A} = \begin{array}{|c|c|c|} \hline * & \mathbf{A}_{\alpha\beta} & * \\ \hline \mathbf{A}_{\beta\alpha} & * & \\ \hline * & * & \mathbf{A}_{\gamma\delta} \\ \hline & \mathbf{A}_{\delta\gamma} & * \\ \hline \end{array}$$

First, observe that

$$\mathbf{A} - \mathbf{A}^{(1)} = \begin{array}{|c|c|c|} \hline * & \mathbf{A}_{\alpha\beta} & \mathbf{0} \\ \hline \mathbf{A}_{\beta\alpha} & * & \\ \hline \mathbf{0} & * & \mathbf{A}_{\gamma\delta} \\ \hline & \mathbf{A}_{\delta\gamma} & * \\ \hline \end{array}$$

We then define two random test matrices, each of size $N \times r$, via

$$\Omega_1 = \begin{bmatrix} \mathbf{G}_\alpha \\ \mathbf{0} \\ \mathbf{G}_\gamma \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \Omega_2 = \begin{bmatrix} \mathbf{0} \\ \mathbf{G}_\beta \\ \mathbf{0} \\ \mathbf{G}_\delta \end{bmatrix}$$

We compute the sample matrices via

$$\mathbf{Y}_1 = \mathbf{A}\Omega_2 - \mathbf{A}^{(1)}\Omega_2 = \begin{bmatrix} \mathbf{A}_{\alpha\beta}\mathbf{G}_\beta \\ * \\ \mathbf{A}_{\gamma\delta}\mathbf{G}_\delta \\ * \end{bmatrix} \quad \text{and} \quad \mathbf{Y}_2 = \mathbf{A}\Omega_1 - \mathbf{A}^{(1)}\Omega_1 = \begin{bmatrix} * \\ \mathbf{A}_{\beta\alpha}\mathbf{G}_\alpha \\ * \\ \mathbf{A}_{\delta\gamma}\mathbf{G}_\gamma \end{bmatrix}. \quad (10)$$

In evaluating \mathbf{Y}_1 and \mathbf{Y}_2 , we use the black-box multiplier to form $\mathbf{A}\Omega_2$ and $\mathbf{A}\Omega_1$, and the compressed representation of $\mathbf{A}^{(1)}$ obtained on the previous level to evaluate $\mathbf{A}^{(1)}\Omega_2$ and $\mathbf{A}^{(1)}\Omega_1$. We now obtain orthonormal bases for the column spaces of the four sibling interaction matrices by orthonormalizing the pertinent blocks of \mathbf{Y}_1 and \mathbf{Y}_2 :

$$\mathcal{U}_\alpha = \text{qr}(\mathbf{Y}_1(I_\alpha, :)), \quad \mathcal{U}_\beta = \text{qr}(\mathbf{Y}_2(I_\beta, :)), \quad \mathcal{U}_\gamma = \text{qr}(\mathbf{Y}_1(I_\gamma, :)), \quad \mathcal{U}_\delta = \text{qr}(\mathbf{Y}_2(I_\delta, :)).$$

It remains to construct the ON-bases for the corresponding row spaces and the compressed sibling interaction matrices. To this end, we form two new test matrices, both of size $N \times r$, via

$$\Omega_1 = \begin{bmatrix} \mathcal{U}_\alpha \\ \mathbf{0} \\ \mathcal{U}_\gamma \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \Omega_2 = \begin{bmatrix} \mathbf{0} \\ \mathcal{U}_\beta \\ \mathbf{0} \\ \mathcal{U}_\delta \end{bmatrix}.$$

Then the sample matrices are computed via

$$\mathbf{Z}_1 = \mathbf{A}^*\Omega_2 - (\mathbf{A}^{(1)})^*\Omega_2 = \begin{bmatrix} \mathbf{A}_{\beta\alpha}^*\mathcal{U}_\beta \\ * \\ \mathbf{A}_{\delta\gamma}^*\mathcal{U}_\delta \\ * \end{bmatrix} \quad \text{and} \quad \mathbf{Z}_2 = \mathbf{A}^*\Omega_1 - (\mathbf{A}^{(1)})^*\Omega_1 = \begin{bmatrix} * \\ \mathbf{A}_{\alpha\beta}^*\mathcal{U}_\alpha \\ * \\ \mathbf{A}_{\gamma\delta}^*\mathcal{U}_\gamma \end{bmatrix}.$$

We obtain diagonal compressed sibling interaction matrices by taking a sequence of dense SVDs of the relevant sub-blocks, cf. (8),

$$\begin{aligned} [\mathcal{V}_\alpha, \mathbf{B}_{\beta\alpha}, \hat{\mathbf{U}}_\beta] &= \text{svd}(\mathbf{Z}_1(I_\alpha, :)), \\ [\mathcal{V}_\beta, \mathbf{B}_{\alpha\beta}, \hat{\mathbf{U}}_\alpha] &= \text{svd}(\mathbf{Z}_2(I_\beta, :)), \\ [\mathcal{V}_\gamma, \mathbf{B}_{\delta\gamma}, \hat{\mathbf{U}}_\delta] &= \text{svd}(\mathbf{Z}_1(I_\gamma, :)), \\ [\mathcal{V}_\delta, \mathbf{B}_{\gamma\delta}, \hat{\mathbf{U}}_\gamma] &= \text{svd}(\mathbf{Z}_2(I_\delta, :)). \end{aligned}$$

Finally update the bases for the column-spaces, cf (9),

$$\mathcal{U}_\alpha \leftarrow \mathcal{U}_\alpha \hat{\mathbf{U}}_\alpha, \quad \mathcal{U}_\beta \leftarrow \mathcal{U}_\beta \hat{\mathbf{U}}_\beta, \quad \mathcal{U}_\gamma \leftarrow \mathcal{U}_\gamma \hat{\mathbf{U}}_\gamma, \quad \mathcal{U}_\delta \leftarrow \mathcal{U}_\delta \hat{\mathbf{U}}_\delta.$$

Processing levels 2 through $L - 1$: The processing of every level proceeds in a manner completely analogous to the processing of level 1. The relevant formulas are given in Figure 4.

Processing the leaves: Once all L levels have been traversed using the procedure described, compressed representations of all off-diagonal blocks will have been computed. At this point, all that remains is to extract the

diagonal blocks. We illustrate the process for a simplistic example of a tree with only $L = 2$ levels (beyond the root). Letting $\{\alpha, \beta, \gamma, \delta\}$ denote the leaf nodes, our task is then to extract the blocks marked in red:

$$\mathbf{A} = \begin{array}{|c|c|c|c|} \hline & \mathbf{D}_\alpha & * & * \\ \hline * & \mathbf{D}_\beta & & \\ \hline & \mathbf{D}_\gamma & * & \\ \hline * & * & \mathbf{D}_\delta & \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline & \mathbf{D}_\alpha & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{D}_\beta & & \\ \hline & \mathbf{0} & \mathbf{D}_\gamma & \mathbf{0} \\ \hline & \mathbf{0} & \mathbf{0} & \mathbf{D}_\delta \\ \hline \end{array} + \mathbf{A}^{(2)}.$$

Since the diagonal blocks are not rank-deficient, we will extract them directly, without using randomized sampling. To describe the process, we assume at first (for simplicity) that every diagonal block has the same size, $m \times m$. We then choose a test matrix of size $N \times m$

$$\Omega = \begin{bmatrix} \mathbf{I}_m \\ \mathbf{I}_m \\ \mathbf{I}_m \\ \mathbf{I}_m \end{bmatrix},$$

and trivially extract the diagonal blocks via the sampling

$$\mathbf{Y} = \mathbf{A}\Omega - \mathbf{A}^{(2)}\Omega = \begin{bmatrix} \mathbf{D}_\alpha \\ \mathbf{D}_\beta \\ \mathbf{D}_\gamma \\ \mathbf{D}_\delta \end{bmatrix}.$$

The diagonal blocks can then be read off directly from \mathbf{Y} .

For the general case where the leaves may be of different sizes, we form a test matrix Ω of size $N \times m$, where $m = \max\{n_\tau : \tau \text{ is a leaf}\}$, such that

$$\Omega(I_\tau, :) = [\mathbf{I}_{n_\tau} \text{zeros}(n_\tau, m - n_\tau)].$$

In other words, we simply pad a few zero columns at the end.

The entire algorithm is summarized in Figure 4.

3.1. Asymptotic complexity. Let L denote the number of levels in the tree. We find that $L \sim \log N$. Let T_{flop} denote the time required for a flop, let T_{mult} denote the time required to apply \mathbf{A} or \mathbf{A}^* to a vector, and let $T_{\mathbf{A}^{(\ell)}}$ denote the time required to apply $\mathbf{A}^{(\ell)}$ to a vector. Then the cost of processing level ℓ is

$$T_\ell \sim T_{\text{mult}} \times k + T_{\text{flop}} \times 2^\ell k \frac{N}{2^\ell} + T_{\mathbf{A}^{(\ell)}} \times k, \quad (11)$$

since on level ℓ there are 2^ℓ blocks to be processed, and each “long” matrix at this level has height $2^{-\ell} N$ and width $O(k)$. Further, we find that the cost of applying $\mathbf{A}^{(\ell)}$ to a single vector is

$$T_{\mathbf{A}^{(\ell)}} \sim T_{\text{flop}} \times \sum_{j=0}^{\ell} 2^j k \frac{N}{2^j} \sim T_{\text{flop}} \times \ell k N. \quad (12)$$

Combining (11) and (12) and summing from $\ell = 0$ to $\ell = L$, we find (using that $L \sim \log N$)

$$T_{\text{compress}} \sim T_{\text{mult}} \times k \log N + T_{\text{flop}} \times k^2 N (\log N)^2. \quad (13)$$

Build compressed representations of all off-diagonal blocks.

loop over levels $\ell = 0 : (L - 1)$

Build the random matrices Ω_1 and Ω_2 .

$$\Omega_1 = \text{zeros}(n, r)$$

$$\Omega_2 = \text{zeros}(n, r)$$

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of box τ .

$$\Omega_1(I_\alpha, :) = \text{randn}(n_\alpha, r)$$

$$\Omega_2(I_\beta, :) = \text{randn}(n_\beta, r)$$

end loop

Apply \mathbf{A} to build the samples for the incoming basis matrices.

$$\mathbf{Y}_1 = \mathbf{A}\Omega_2 - \mathbf{A}^{(\ell)}\Omega_2$$

$$\mathbf{Y}_2 = \mathbf{A}\Omega_1 - \mathbf{A}^{(\ell)}\Omega_1$$

Orthonormalize the sample matrices to build the incoming basis matrices.

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of box τ .

$$\mathcal{U}_\alpha = \text{qr}(\mathbf{Y}_1(I_\alpha, :)).$$

$$\mathcal{U}_\beta = \text{qr}(\mathbf{Y}_2(I_\beta, :)).$$

$$\Omega_1(I_\alpha, :) = \mathcal{U}_\alpha$$

$$\Omega_2(I_\beta, :) = \mathcal{U}_\beta$$

end loop

Apply \mathbf{A}^ to build the samples for the outgoing basis matrices.*

$$\mathbf{Z}_1 = \mathbf{A}^*\Omega_2 - (\mathbf{A}^{(\ell)})^*\Omega_2$$

$$\mathbf{Z}_2 = \mathbf{A}^*\Omega_1 - (\mathbf{A}^{(\ell)})^*\Omega_1$$

Take local SVDs to build incoming basis matrices and sibling interaction matrices.

We determine the actual rank, and update the \mathcal{U}_ basis matrices accordingly.*

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of box τ .

$$[\mathcal{V}_\alpha, \mathbf{B}_{\beta\alpha}, \hat{\mathbf{U}}_\beta] = \text{svd}(\mathbf{Z}_1(I_\alpha, :), \varepsilon).$$

$$[\mathcal{V}_\beta, \mathbf{B}_{\alpha\beta}, \hat{\mathbf{U}}_\alpha] = \text{svd}(\mathbf{Z}_2(I_\beta, :), \varepsilon).$$

$$\mathcal{U}_\beta \leftarrow \mathcal{U}_\beta \hat{\mathbf{U}}_\beta.$$

$$\mathcal{U}_\alpha \leftarrow \mathcal{U}_\alpha \hat{\mathbf{U}}_\alpha.$$

end loop

end loop

Extract the diagonal matrices.

$$n_{\max} = \max\{n_\tau : \tau \text{ is a leaf}\}$$

$$\Omega = \text{zeros}(N, n_{\max})$$

loop over leaf boxes τ

$$\Omega(I_\tau, 1 : n_\tau) = \text{eye}(n_\tau).$$

end loop

$$\mathbf{Y} = \mathbf{A}\Omega - \mathbf{A}^{(L)}\Omega$$

loop over leaf boxes τ

$$\mathbf{D}_\tau = \mathbf{Y}(I_\tau, 1 : n_\tau).$$

end loop

FIGURE 4. Randomized compression of a HODLR matrix.

4. AN ALGORITHM FOR COMPRESSING AN HBS MATRIX

The algorithm for computing a compressed representation of an HBS matrix is a slight variation of the algorithm for a HODLR matrix described in Section 3. The key difference is that the long basis matrices \mathcal{U}_α and \mathcal{V}_α associated with any node now must satisfy a more rigorous requirement. We will accomplish this objective by constructing for every node α , two new “long” sampling matrices \mathcal{Y}_α and \mathcal{Z}_α , each of size $n_\alpha \times r$ that help transmit information from the higher levels to the node α . The presentation will start in Section 4.1 with a description of the modification to the scheme of 3 required to enforce the more rigorous requirement. In this initial description, we will assume that all four “long” matrices associated with a node $(\mathcal{U}_\tau, \mathcal{V}_\tau, \mathcal{Y}_\tau, \mathcal{Z}_\tau)$ are stored explicitly, resulting in an $O(k N \log N)$ memory requirement, just like for the HODLR algorithm. In Section 4.2 we show that while these “long” matrices do need to be temporarily built and processed, they can be stored *implicitly*, which will allow the algorithm to use only $O(k N)$ memory. Finally, Section 4.3 will describe how to construct a representation using interpolatory decompositions in all low-rank approximations.

4.1. A basic scheme for compressing an HBS matrix. Throughout this section, let α denote a node with a parent τ that is not the root, and with a sibling β . We will first describe the process for building the long basis matrices $\{\mathcal{U}_\tau\}_\tau$. To this end, recall that the difference in requirements on the long basis matrices in the two frameworks is as follows:

HODLR framework: The columns of \mathcal{U}_α need to span the columns of $\mathbf{A}(I_\alpha, I_\beta)$

HBS framework: The columns of \mathcal{U}_α need to span the columns of $\mathbf{A}(I_\alpha, I_\alpha^c)$.

The assertion that the requirements on a basis in the HBS framework is more stringent follows from the fact that $I_\beta \subseteq I_\alpha^c$, and that I_β is in general much smaller than I_α^c . Now note that

$$\mathbf{A}(I_\alpha, I_\alpha^c) = [\mathbf{A}(I_\alpha, I_\beta) \mathbf{A}(I_\alpha, I_\tau^c)] \mathbf{P}. \quad (14)$$

In (14), the matrix \mathbf{P} is a permutation matrix whose effect is to reorder the columns. For purposes of constructing a basis for the column space, the matrix \mathbf{P} can be ignored. The idea is now to introduce a new sampling matrix \mathcal{Y}_α of size $n_\alpha \times r$ that encodes all the information that needs to be transmitted from the parent τ . Specifically, we ask that:

The columns of \mathcal{Y}_α span the columns of $\mathbf{A}(I_\alpha, I_\tau^c)$ (to within precision ε).

Then, when processing box α , we will sample $\mathbf{A}(I_\alpha, I_\beta)$ using a Gaussian matrix \mathbf{G}_β of size $n_\beta \times r$ just as in the HODLR algorithm. In the end, we will build \mathcal{U}_α by combining the two sets of samples

$$[\mathcal{U}_\alpha, \mathbf{D}_\alpha, \sim] = \text{svd}([\mathbf{A}(I_\alpha, I_\beta) \mathbf{G}_\beta, \mathcal{Y}_\alpha], r).$$

In other words, we take a matrix $[\mathbf{A}(I_\alpha, I_\beta) \mathbf{G}_\beta, \mathcal{Y}_\alpha]$ of size $n_\alpha \times 2r$ and extract its leading r singular components (the trailing r components are discarded). All that remains is to build the sample matrices \mathcal{Y}_γ and \mathcal{Y}_δ that transmit information to the children $\{\gamma, \delta\}$ of α . To be precise, let J_γ and J_δ denote the local *relative* index vectors, so that

$$I_\gamma = I_\alpha(J_\gamma) \quad \text{and} \quad I_\delta = I_\alpha(J_\delta).$$

Then set

$$\mathcal{Y}_\gamma = \mathcal{U}(J_\gamma, :) \mathbf{D}_\alpha \quad \text{and} \quad \mathcal{Y}_\delta = \mathcal{U}(J_\delta, :) \mathbf{D}_\alpha.$$

The process for building the long basis matrices $\{\mathcal{V}_\alpha\}_\alpha$ is entirely analogous to the process described for building the $\{\mathcal{U}_\alpha\}_\alpha$ matrices. We first recall that the difference between the HODLR and the HBS frameworks is as follows:

HODLR framework: The columns of \mathcal{V}_α need to span the rows of $\mathbf{A}(I_\beta, I_\alpha)$

HBS framework: The columns of \mathcal{V}_α need to span the rows of $\mathbf{A}(I_\alpha^c, I_\alpha)$.

```

loop over levels  $\ell = 0 : (L - 1)$ 
  Build the random matrices  $\Omega_1$  and  $\Omega_2$ .
   $\Omega_1 = \text{zeros}(n, r)$ 
   $\Omega_2 = \text{zeros}(n, r)$ 
  loop over boxes  $\tau$  on level  $\ell$ 
    Let  $\{\alpha, \beta\}$  denote the children of box  $\tau$ .
     $\Omega_1(I_\alpha, :) = \text{randn}(n_\alpha, r)$ 
     $\Omega_2(I_\beta, :) = \text{randn}(n_\beta, r)$ 
  end loop
  Apply  $\mathbf{A}$  to build the samples for the incoming basis matrices.
   $\mathbf{Y}_1 = \mathbf{A}\Omega_2 - \mathbf{A}^{(\ell)}\Omega_2$ 
   $\mathbf{Y}_2 = \mathbf{A}\Omega_1 - \mathbf{A}^{(\ell)}\Omega_1$ 
  Orthonormalize the sample matrices to build the incoming basis matrices.
  loop over boxes  $\tau$  on level  $\ell$ 
    Let  $\{\alpha, \beta\}$  denote the children of box  $\tau$ .
    if ( $\tau$  is the root)
       $\mathbf{Y}_1^{\text{loc}} = \mathbf{Y}_1(I_\alpha, :)$ 
       $\mathbf{Y}_2^{\text{loc}} = \mathbf{Y}_2(I_\beta, :)$ 
    else
       $\mathbf{Y}_1^{\text{loc}} = [\mathbf{Y}_1(I_\alpha, :), \mathbf{Y}_\tau(J_\alpha, :)]$ 
       $\mathbf{Y}_2^{\text{loc}} = [\mathbf{Y}_2(I_\beta, :), \mathbf{Y}_\tau(J_\beta, :)]$ 
    end if
     $[\mathcal{U}_\alpha, \mathbf{y}_\alpha, \sim] = \text{svd}(\mathbf{Y}_1^{\text{loc}}, r)$ .
     $[\mathcal{U}_\beta, \mathbf{y}_\beta, \sim] = \text{svd}(\mathbf{Y}_2^{\text{loc}}, r)$ .
     $\mathcal{Y}_\alpha = \mathcal{U}_\alpha \text{diag}(\mathbf{y}_\alpha)$ .
     $\mathcal{Y}_\beta = \mathcal{U}_\beta \text{diag}(\mathbf{y}_\beta)$ .
     $\Omega_1(I_\alpha, :) = \mathcal{U}_\alpha$ 
     $\Omega_2(I_\beta, :) = \mathcal{U}_\beta$ 
     $\mathbf{U}_\tau = \begin{bmatrix} \mathcal{U}_\alpha^* \mathcal{U}_\tau(J_\alpha, :) \\ \mathcal{U}_\beta^* \mathcal{U}_\tau(J_\beta, :) \end{bmatrix}$ 
  end loop

```

Apply \mathbf{A}^ to build the samples for the outgoing basis matrices.*

$$\mathbf{Z}_1 = \mathbf{A}^* \Omega_2 - (\mathbf{A}^{(\ell)})^* \Omega_2$$

$$\mathbf{Z}_2 = \mathbf{A}^* \Omega_1 - (\mathbf{A}^{(\ell)})^* \Omega_1$$

Take local SVDs to build incoming basis matrices and sibling interaction matrices.

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of box τ .

if (τ is the root)

$$\mathbf{Z}_1^{\text{loc}} = \mathbf{Z}_1(I_\alpha, :)$$

$$\mathbf{Z}_2^{\text{loc}} = \mathbf{Z}_2(I_\beta, :)$$

else

$$\mathbf{Z}_1^{\text{loc}} = [\mathbf{Z}_1(I_\alpha, :), \mathbf{Z}_\tau(J_\alpha, :)]$$

$$\mathbf{Z}_2^{\text{loc}} = [\mathbf{Z}_2(I_\beta, :), \mathbf{Z}_\tau(J_\beta, :)]$$

end if

$$[\mathcal{V}_\alpha, \mathbf{b}_{21}, \mathbf{X}_1] = \text{svd}(\mathbf{Z}_1^{\text{loc}}, r)$$

$$[\mathcal{V}_\beta, \mathbf{b}_{12}, \mathbf{X}_2] = \text{svd}(\mathbf{Z}_2^{\text{loc}}, r)$$

$$\mathcal{Z}_\alpha = \mathcal{V}_\alpha \text{diag}(\mathbf{b}_{21})$$

$$\mathcal{Z}_\beta = \mathcal{V}_\beta \text{diag}(\mathbf{b}_{12})$$

$$\mathbf{B}_{\alpha\beta} = \mathbf{X}_1(1:r,:) \text{diag}(\mathbf{b}_{12})$$

$$\mathbf{B}_{\beta\alpha} = \mathbf{X}_2(1:r,:) \text{diag}(\mathbf{b}_{21})$$

$$\mathbf{V}_\tau = \begin{bmatrix} \mathcal{V}_\alpha^* \mathcal{V}_\tau(J_\alpha, :) \\ \mathcal{V}_\beta^* \mathcal{V}_\tau(J_\beta, :) \end{bmatrix}$$

end loop

end loop

Extract the diagonal matrices.

$$n_{\max} = \max \{n_\tau : \tau \text{ is a leaf}\}$$

$$\Omega = \text{zeros}(N, n_{\max})$$

loop over leaf boxes τ

$$\Omega(I_\tau, 1:n_\tau) = \text{eye}(n_\tau)$$

end loop

$$\mathbf{Y} = \mathbf{A}\Omega - \mathbf{A}^{(L)}\Omega$$

loop over leaf boxes τ

$$\mathbf{D}_\tau = \mathbf{Y}(I_\tau, 1:n_\tau)$$

end loop

FIGURE 5. A basic scheme for compressing an HBS matrix.

With \mathbf{P} again denoting a permutation matrix, we have

$$\mathbf{A}(I_\alpha^c, I_\alpha) = \mathbf{P} \begin{bmatrix} \mathbf{A}(I_\beta, I_\alpha) \\ \mathbf{A}(I_\tau^c, I_\alpha) \end{bmatrix}.$$

It follows that the role that was played by \mathcal{Y}_α in the construction of \mathcal{U}_α is now played by a sampling matrix \mathcal{Z}_α of size $n_\alpha \times r$ that satisfies

The columns of \mathcal{Z}_α span the rows of $\mathbf{A}(I_\tau^c, I_\alpha)$.

The algorithm for computing the HBS representation of a matrix is now given in Figure 5.

4.2. A storage efficient scheme for compressing an HBS matrix. The scheme described in Section 4.1 assumes that all “long” basis and spanning matrices (\mathcal{U}_τ , \mathcal{V}_τ , \mathcal{Y}_τ , \mathcal{Z}_τ) are stored explicitly, resulting in an $O(k N \log N)$ storage requirement. We will now demonstrate that in fact, only $O(k N)$ is required.

First, observe that in the HBS framework, we only need to keep at hand the long basis matrices \mathcal{U}_τ and \mathcal{V}_τ for nodes τ on the level ℓ that is currently being processed. In the HODLR compression algorithm in Figure 4, we needed the long basis matrices associated with nodes on coarser levels in order to apply $\mathbf{A}^{(\ell)}$, but in the HBS framework, all we need in order to apply $\mathbf{A}^{(\ell)}$ is the long basis matrices $\{\mathcal{U}_\tau, \mathcal{V}_\tau\}$ on the level currently being processed, and then only the short basis matrices \mathbf{U}_τ and \mathbf{V}_τ for any box τ on a level coarser than ℓ , cf. the algorithm in Figure 15.

Next, observe that the long “spanning” matrices \mathcal{Y}_τ and \mathcal{Z}_τ that were introduced in Section 4.1 do not need to be stored explicitly either. The reason is that these matrices can be expressed in terms of the long basis matrices \mathcal{U}_τ and \mathcal{V}_τ . In fact, in the algorithm in Figure 4, we compute \mathcal{Y}_τ and \mathcal{Z}_τ via the relations

$$\mathcal{Y}_\tau = \mathcal{U}_\tau \operatorname{diag}(\mathbf{y}_\tau) \quad \text{and} \quad \mathcal{Z}_\tau = \mathcal{V}_\tau \operatorname{diag}(\mathbf{z}_\tau).$$

Since the long basis matrices \mathcal{U}_τ and \mathcal{V}_τ are available during the processing of level ℓ , we only need to store the short vectors \mathbf{y}_τ and \mathbf{z}_τ , and can then construct \mathcal{Y}_τ and \mathcal{Z}_τ when they are actually needed.

The memory efficient algorithm resulting from exploiting the observations described in this section is summarized in Figure 6.

4.3. Adaptive rank determination and conversion to the HBS-ID format. The schemes presented in Sections 4.1 and 4.2 do not adaptively determine the ranks of the off-diagonal blocks being compressed. Instead, every block is factored using a preset uniform rank ℓ that must be picked to be larger than any actual numerical rank encountered. It is possible to incorporate adaptive rank determination into the scheme, but we found it easier to perform this step in a second sweep that travels through the tree in the opposite direction: from smaller boxes to larger. In this second sweep, we also convert the standard HBS format to the HBS-ID format, which leads to a slight improvement in storage requirements, and improves interpretability of the sibling interaction matrices, as discussed in Section 2.9 and Definition 3.

The conversion to the HBS-ID is a “post-processing” step, so in what follows, we assume that the compression algorithm in Figure 6 has already been executed so that both the HBS basis matrices \mathcal{U}_τ and \mathcal{V}_τ , and the sample matrix \mathcal{Y}_τ and \mathcal{Z}_τ are available for every node (these are stored *implicitly* in terms of the short basis matrices \mathbf{U}_τ and \mathbf{V}_τ , as described in Section 4.2).

The first step is to sweep over all leaves τ in the tree. For each leaf, we now have available spanning matrices \mathcal{Y}_τ and \mathcal{Z}_τ whose columns span the columns of $\mathbf{A}(I_\tau, I_\tau^c)$ and $\mathbf{A}(I_\tau^c, I_\tau)^*$, respectively. In order to find a set of spanning rows $\tilde{I}_\tau^{\text{in}}$ of $\mathbf{A}(I_\tau, I_\tau^c)$ and a set of spanning columns $\tilde{I}_\tau^{\text{out}}$ of $\mathbf{A}(I_\tau^c, I_\tau)^*$, all we need to do is to compute interpolatory decompositions (IDs) of the small matrices \mathcal{Y}_τ and \mathcal{Z}_τ , cf. Remark 4,

$$[\mathbf{T}_{\text{in}}, J_{\text{in}}] = \text{id}(\mathcal{Y}_\tau^*, \varepsilon), \quad \text{and} \quad [\mathbf{T}_{\text{out}}, J_{\text{out}}] = \text{id}(\mathcal{Z}_\tau^*, \varepsilon). \quad (15)$$

In equation (15), we give the computational tolerance ε as an input parameter. This reveals the “true” ε -ranks k_{in} and k_{out} . Then the skeleton index vectors for τ are given by

$$\tilde{I}_\tau^{\text{in}} = I_\tau(J_{\text{in}}(1 : k_{\text{in}})), \quad \text{and} \quad \tilde{I}_\tau^{\text{out}} = I_\tau(J_{\text{out}}(1 : k_{\text{out}})).$$

Now define the *subsampled* basis matrices $\mathbf{U}_\tau^{\text{samp}}$ and $\mathbf{V}_\tau^{\text{samp}}$ via

$$\mathbf{U}_\tau^{\text{samp}} = \mathbf{U}(J_{\text{in}}(1 : k_{\text{in}}), :), \quad \text{and} \quad \mathbf{V}_\tau^{\text{samp}} = \mathbf{V}(J_{\text{out}}(1 : k_{\text{out}}), :). \quad (16)$$

Once all leaves have been processed in this manner, we can determine the sibling interaction matrices in the HBS-ID representation. Let $\{\alpha, \beta\}$ denote a sibling pair consisting of two leaves. First observe that, by definition,

$$\mathbf{B}_{\alpha, \beta}^{\text{skel}} = \mathbf{A}(\tilde{I}_\alpha^{\text{in}}, \tilde{I}_\beta^{\text{out}}). \quad (17)$$

Next, recall that

$$\mathbf{A}(I_\alpha, I_\beta) = \mathbf{U}_\alpha \mathbf{B}_{\alpha, \beta} \mathbf{V}_\beta^*. \quad (18)$$

```

loop over levels  $\ell = 0 : (L - 1)$ 
  Build the random matrices  $\Omega_1$  and  $\Omega_2$ .
   $\Omega_1 = \text{zeros}(n, r)$ 
   $\Omega_2 = \text{zeros}(n, r)$ 
  loop over boxes  $\tau$  on level  $\ell$ 
    Let  $\{\alpha, \beta\}$  denote the children of box  $\tau$ .
     $\Omega_1(I_\alpha, :) = \text{randn}(n_\alpha, r)$ 
     $\Omega_2(I_\beta, :) = \text{randn}(n_\beta, r)$ 
  end loop

  Apply  $\mathbf{A}$  to build the samples for the incoming basis matrices.
   $\mathbf{Y}_1 = \mathbf{A}\Omega_2 - \mathbf{A}^{(\ell)}\Omega_2$ 
   $\mathbf{Y}_2 = \mathbf{A}\Omega_1 - \mathbf{A}^{(\ell)}\Omega_1$ 

  Orthonormalize the sample matrices to build the incoming basis matrices.
  loop over boxes  $\tau$  on level  $\ell$ 
    Let  $\{\alpha, \beta\}$  denote the children of box  $\tau$ .
    if ( $\tau$  is the root)
       $\mathbf{Y}_1^{\text{loc}} = \mathbf{Y}_1(I_\alpha, :)$ 
       $\mathbf{Y}_2^{\text{loc}} = \mathbf{Y}_2(I_\beta, :)$ 
    else
       $\mathbf{Y}_1^{\text{loc}} = [\mathbf{Y}_1(I_\alpha, :), \mathcal{U}_\tau(J_\alpha, :) \text{diag}(\mathbf{y}_\tau)]$ 
       $\mathbf{Y}_2^{\text{loc}} = [\mathbf{Y}_2(I_\beta, :), \mathcal{U}_\tau(J_\beta, :) \text{diag}(\mathbf{y}_\tau)]$ 
    end if
     $[\mathcal{U}_\alpha, \mathbf{y}_\alpha, \sim] = \text{svd}(\mathbf{Y}_1^{\text{loc}}, r).$ 
     $[\mathcal{U}_\beta, \mathbf{y}_\beta, \sim] = \text{svd}(\mathbf{Y}_2^{\text{loc}}, r).$ 
     $\Omega_1(I_\alpha, :) = \mathcal{U}_\alpha$ 
     $\Omega_2(I_\beta, :) = \mathcal{U}_\beta$ 
     $\mathbf{U}_\tau = \begin{bmatrix} \mathcal{U}_\alpha^* \mathcal{U}_\tau(J_\alpha, :) \\ \mathcal{U}_\beta^* \mathcal{U}_\tau(J_\beta, :) \end{bmatrix}.$ 
    Delete  $\mathcal{U}_\tau$ .
  end loop

```

Apply \mathbf{A}^ to build the samples for the outgoing basis matrices.*

$$\mathbf{Z}_1 = \mathbf{A}^* \Omega_2 - (\mathbf{A}^{(\ell)})^* \Omega_2$$

$$\mathbf{Z}_2 = \mathbf{A}^* \Omega_1 - (\mathbf{A}^{(\ell)})^* \Omega_1$$

Take local SVDs to build incoming basis matrices and sibling interaction matrices.

loop over boxes τ on level ℓ
 Let $\{\alpha, \beta\}$ denote the children of box τ .

if (τ is the root)

$$\mathbf{Z}_1^{\text{loc}} = \mathbf{Z}_1(I_\alpha, :)$$

$$\mathbf{Z}_2^{\text{loc}} = \mathbf{Z}_2(I_\beta, :)$$

else

$$\mathbf{Z}_1^{\text{loc}} = [\mathbf{Z}_1(I_\alpha, :), \mathcal{V}_\tau(J_\alpha, :) \text{diag}(\mathbf{z}_\tau)]$$

$$\mathbf{Z}_2^{\text{loc}} = [\mathbf{Z}_2(I_\beta, :), \mathcal{V}_\tau(J_\beta, :) \text{diag}(\mathbf{z}_\tau)]$$

end if

$$[\mathcal{V}_\alpha, \mathbf{b}_{21}, \mathbf{X}_1] = \text{svd}(\mathbf{Z}_1^{\text{loc}}, r).$$

$$[\mathcal{V}_\beta, \mathbf{b}_{12}, \mathbf{X}_2] = \text{svd}(\mathbf{Z}_2^{\text{loc}}, r).$$

$$\mathbf{z}_\alpha = \mathbf{b}_{21}$$

$$\mathbf{z}_\beta = \mathbf{b}_{12}$$

$$\mathbf{B}_{\alpha\beta} = \mathbf{X}_1(1:r, :) \text{diag}(\mathbf{b}_{12})$$

$$\mathbf{B}_{\beta\alpha} = \mathbf{X}_2(1:r, :) \text{diag}(\mathbf{b}_{21})$$

$$\mathbf{V}_\tau = \begin{bmatrix} \mathcal{V}_\alpha^* \mathcal{V}_\tau(J_\alpha, :) \\ \mathcal{V}_\beta^* \mathcal{V}_\tau(J_\beta, :) \end{bmatrix}.$$

Delete \mathcal{V}_τ .

end loop

end loop

Extract the diagonal matrices.

$$n_{\max} = \max \{n_\tau : \tau \text{ is a leaf}\}$$

$$\Omega = \text{zeros}(N, n_{\max})$$

loop over leaf boxes τ

$$\Omega(I_\tau, 1:n_\tau) = \text{eye}(n_\tau).$$

end loop

$$\mathbf{Y} = \mathbf{A}\Omega - \mathbf{A}^{(L)}\Omega$$

loop over leaf boxes τ

$$\mathbf{D}_\tau = \mathbf{Y}(I_\tau, 1:n_\tau).$$

end loop

FIGURE 6. A storage efficient algorithm for compressing an HBS matrix.

Combining (16), (17), and (18), we find that

$$\mathbf{B}_{\alpha, \beta}^{\text{skel}} = \mathbf{U}_\alpha^{\text{samp}} \mathbf{B}_{\alpha, \beta} (\mathbf{V}_\beta^{\text{samp}})^*.$$

Once all leaves have been processed, we next proceed to the parent nodes. We do this by transversing the tree in the other direction, going from smaller to larger boxes. When a box τ is processed, its children $\{\alpha, \beta\}$ have already been processed. The key observation is now that if we set $\hat{I}_\tau^{\text{in}} = \tilde{I}_\alpha^{\text{in}} \cup \tilde{I}_\beta^{\text{in}}$ and $\hat{I}_\tau^{\text{out}} = \tilde{I}_\alpha^{\text{out}} \cup \tilde{I}_\beta^{\text{out}}$, then these index vectors form skeletons for τ . These skeletons are inefficient, but by simply compressing the corresponding rows of \mathbf{Y}_τ and \mathbf{Z}_τ , we can build the skeletons and the interpolation matrices associated with τ . Due to the self-similarity between levels in the HBS representation, the compression is entirely analogous to the compression of a leaf, with the only difference that the role played by the basis matrices \mathcal{U}_τ and \mathcal{V}_τ for a leaf, are now played by the sub-sampled matrices \mathbf{U}_{tmp} and \mathbf{V}_{tmp} which represent the restriction of \mathcal{U}_τ and

```

Execute the compression algorithm described in Figure 6.

loop over leaves  $\tau$ 
   $\mathbf{Y}_{\text{tmp}} = \mathbf{U}_\tau \text{diag}(\mathbf{y}_\tau)$ 
   $\mathbf{Z}_{\text{tmp}} = \mathbf{V}_\tau \text{diag}(\mathbf{z}_\tau)$ 
   $[\mathbf{T}_{\text{in}}, J_{\text{in}}] = \text{id}(\mathbf{Y}_{\text{tmp}}^*, \varepsilon)$ 
   $[\mathbf{T}_{\text{out}}, J_{\text{out}}] = \text{id}(\mathbf{Z}_{\text{tmp}}^*, \varepsilon)$ 
   $\mathbf{U}_\tau^{\text{samp}} = \mathbf{U}_\tau(J_{\text{in}}(1:k_{\text{in}}), :)$ .
   $\mathbf{V}_\tau^{\text{samp}} = \mathbf{V}_\tau(J_{\text{out}}(1:k_{\text{out}}), :)$ .
end loop

```

```

loop over levels  $\ell = (L-1) : (-1) : 1$ 
  loop over boxes  $\tau$  on level  $\ell$ 
    Let  $\{\alpha, \beta\}$  denote the children of box  $\tau$ .
     $\mathbf{U}_{\text{tmp}} = \begin{bmatrix} \mathbf{U}_\alpha^{\text{samp}} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_\beta^{\text{samp}} \end{bmatrix} \mathbf{U}_\tau$ .
     $\mathbf{V}_{\text{tmp}} = \begin{bmatrix} \mathbf{V}_\alpha^{\text{samp}} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_\beta^{\text{samp}} \end{bmatrix} \mathbf{V}_\tau$ .
     $\mathbf{Y}_{\text{tmp}} = \mathbf{U}_{\text{tmp}} \text{diag}(\mathbf{y}_\tau)$ 
     $\mathbf{Z}_{\text{tmp}} = \mathbf{V}_{\text{tmp}} \text{diag}(\mathbf{z}_\tau)$ 
     $[\mathbf{T}_{\text{in}}, J_{\text{in}}] = \text{id}(\mathbf{Y}_{\text{tmp}}^*, \varepsilon)$ 
     $[\mathbf{T}_{\text{out}}, J_{\text{out}}] = \text{id}(\mathbf{Z}_{\text{tmp}}^*, \varepsilon)$ 
     $\mathbf{U}_\tau^{\text{samp}} = \mathbf{U}_{\text{tmp}}(J_{\text{in}}(1:k_{\text{in}}), :)$ .
     $\mathbf{V}_\tau^{\text{samp}} = \mathbf{V}_{\text{tmp}}(J_{\text{out}}(1:k_{\text{out}}), :)$ .
     $\mathbf{B}_{\text{skel}}^{\alpha, \beta} = \mathbf{U}_\alpha^{\text{samp}} \mathbf{B}_{\alpha, \beta} (\mathbf{V}_\beta^{\text{samp}})^*$ .
     $\mathbf{B}_{\text{skel}}^{\beta, \alpha} = \mathbf{U}_\beta^{\text{samp}} \mathbf{B}_{\beta, \alpha} (\mathbf{V}_\alpha^{\text{samp}})^*$ .
  end loop
end loop
Let  $\{\alpha, \beta\}$  denote the children of the root.
 $\mathbf{B}_{\text{skel}}^{\alpha, \beta} = \mathbf{U}_\alpha^{\text{samp}} \mathbf{B}_{\alpha, \beta} (\mathbf{V}_\beta^{\text{samp}})^*$ .
 $\mathbf{B}_{\text{skel}}^{\beta, \alpha} = \mathbf{U}_\beta^{\text{samp}} \mathbf{B}_{\beta, \alpha} (\mathbf{V}_\alpha^{\text{samp}})^*$ .

```

FIGURE 7. An algorithm for computing the HBS-ID representation of a given matrix \mathbf{A} . This algorithm adaptively determines the ranks of the off-diagonal blocks of \mathbf{A} .

\mathbf{V}_τ to the index rows and columns indicated by the index vectors \hat{I}_τ^{in} and $\hat{I}_\tau^{\text{out}}$, respectively. The entire process is summarized in Figure 7.

4.4. Asymptotic complexity. The asymptotic complexity for the HBSID algorithm is very similar to that for the HODLR algorithm. The key difference is that since $\mathbf{A}^{(\ell)}$ is now applied using nested basis, we find

$$T_{\mathbf{A}^{(\ell)}} \sim T_{\text{flop}} \times \left(2^\ell k \frac{N}{2^\ell} + \sum_{j=0}^{\ell-1} 2^j k^2 \right) \sim T_{\text{flop}} \times k N.$$

In other words, the asymptotic complexity of applying $T_{\mathbf{A}^{(\ell)}}$ is now less by a factor of $O(\ell)$. In consequence,

$$T_{\text{compress}} \sim T_{\text{mult}} \times k \log N + T_{\text{flop}} \times k^2 N \log N. \quad (19)$$

5. NUMERICAL EXPERIMENTS

In this section, we present results from numerical experiments that substantiate claims on asymptotic complexity made in sections 3.1 and 4.4, and demonstrate that the practical execution time is very competitive (in other words, that the scaling constants suppressed in the asymptotic analysis are moderate). We investigate four different test problems: In Section 5.1 we apply the randomized compression schemes to a discretized boundary integral operator for which other compression techniques are already available. This allows us to benchmark the new algorithms and verify their accuracy. In Section 5.2 we demonstrate how the proposed scheme can be used to form a compressed representation of a product of two compressed matrices, thus demonstrating a way of circumventing the need for a complex and time-consuming structured matrix-matrix multiplication. In Section 5.3 we apply the schemes to a potential evaluation problem where we use the Fast Multipole Method (FMM) as the ‘‘black-box’’ matrix-vector multiplication scheme (note that the data sparse format implicit in the FMM is much more cumbersome to invert than the HODLR and HBSID formats). In Section 5.4 we apply the scheme to compress large dense matrices that arise in the classical ‘‘nested dissection’’ or ‘‘multifrontal’’ direct solvers

for the sparse matrices arising from finite element or finite difference discretization of elliptic PDEs. The ability to efficiently manipulate such matrices allows for the construction of $O(N)$ complexity direct solvers for the associated sparse linear systems.

For each of the four test problems, we compare two different techniques for computing a data-sparse representation of \mathbf{A} : (1) The randomized technique for computing an HODLR-representation in Figure 4. (2) The randomized technique for computing an HBSID-representation in Figure 7. For each technique, we report the following quantities:

N	The number of DOFs (so that \mathbf{A} is of size $N \times N$).
ℓ	The number of random vectors used at each level (ℓ must be larger than the maximal rank).
k	The largest rank encountered in the compression.
N_{matvec}	The number of applications of \mathbf{A} required (so that $N_{\text{matvec}} = (L + 1) \times \ell \sim \log(N) \times \ell$).
T_{compress}	The time required for compression (in seconds).
T_{net}	The time required for compression, excluding time to apply \mathbf{A} and \mathbf{A}^* (in seconds).
T_{app}	The time required for applying the compressed matrix to a vector (in seconds).
M	The amount of memory required to store \mathbf{A} (in MB).

The reason that we report the time T_{net} (that does not count time spent in the black-box matrix-vector multiplier) is to validate our claims (13) and (19) regarding the asymptotic complexity of the method. To summarize, our predictions are, for the HODLR algorithm

$$T_{\text{net}} \sim N (\log N)^2, \quad T_{\text{app}} \sim N \log N, \quad M \sim N \log(N),$$

and for the HBSID algorithm

$$T_{\text{net}} \sim N \log N, \quad T_{\text{app}} \sim N, \quad M \sim N.$$

In addition to the timings, we computed a randomized estimate E of the compression error, computed as follows: We drew ten vectors $\{\omega_i\}_{i=1}^{10}$ of unit length from a uniform distribution on the unit sphere in \mathbb{R}^N . Then E is defined via

$$E = \max_{1 \leq i \leq 10} \frac{\|\mathbf{A}\omega_i - \mathbf{A}_{\text{compressed}}\omega_i\|}{\|\mathbf{A}\omega_i\|}.$$

5.1. Compressing a Boundary Integral Equation. Our first numerical example concerns compression of a discretized version of the Boundary Integral Equation (BIE)

$$\frac{1}{2}q(\mathbf{x}) + \int_{\Gamma} \frac{(\mathbf{x} - \mathbf{y}) \cdot \mathbf{n}(\mathbf{y})}{4\pi|\mathbf{x} - \mathbf{y}|^2} q(\mathbf{y}) ds(\mathbf{y}) = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \quad (20)$$

where Γ is the simple curve shown in Figure 8, and where $\mathbf{n}(\mathbf{y})$ is the outwards pointing unit normal of Γ at \mathbf{y} . The BIE (20) is a standard integral equation formulation of the Laplace equation with boundary condition f on the domain interior to Γ .¹ We discretize the BIE (20) using the Nyström method on N equispaced points on Γ , with the Trapezoidal rule as the quadrature. Note that the kernel in (20) is smooth, so the Trapezoidal rule has exponential convergence. This problem is slightly artificial in that only about 200 points² are needed to attain full double precision accuracy in the discretization. We include it for bench-marking purposes to verify the scaling of the proposed method.

To be precise, the matrix \mathbf{A} used in this numerical experiment is itself an HBS representation of the matrix resulting from discretization of (20), computed using the technique of [28]. To minimize the risk of spurious effects due to both the “exact” and the computed \mathbf{A} being HBS representations, we used a much higher precision in computing the “exact” \mathbf{A} , and also a shifted tree structure to avoid having the compressed blocks of our reference \mathbf{A} align with the compressed blocks constructed by the randomized sampling algorithms.

¹Verify!

²Check.

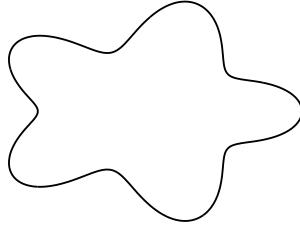


FIGURE 8. Contour Γ on which the BIE (20) in Section 5.1 is defined.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	k
400	3 x 35	0.065	0.035	0.001	0.5	157.7	4.68e-11	28
800	4 x 35	0.107	0.033	0.001	1.0	169.9	3.67e-11	28
1600	5 x 35	0.267	0.081	0.003	2.2	179.1	3.09e-11	28
3200	6 x 35	0.617	0.188	0.006	4.5	186.2	2.57e-11	28
6400	7 x 35	1.458	0.456	0.011	9.5	194.0	1.99e-11	29
12800	8 x 35	3.453	1.198	0.024	19.6	200.5	1.80e-11	30
25600	9 x 35	8.065	2.998	0.047	39.7	203.1	1.74e-11	29
51200	10 x 35	18.362	7.140	0.099	82.1	210.3	1.14e-11	30
102400	11 x 35	41.985	17.366	0.204	166.6	213.2	1.24e-11	30

TABLE 1. Compression to the HODLR format using Algorithm 4 of the double layer integral equation described in Section 5.1. Here $\varepsilon = 10^{-9}$ and $\ell = 35$.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	k
400	3 x 35	0.098	0.068	0.001	0.4	127.7	1.40e-09	25
800	4 x 35	0.163	0.090	0.002	0.7	113.6	1.30e-09	25
1600	5 x 35	0.381	0.193	0.005	1.3	104.3	1.42e-09	24
3200	6 x 35	0.840	0.411	0.009	2.4	98.3	1.32e-09	24
6400	7 x 35	1.881	0.878	0.019	4.6	94.2	1.56e-09	24
12800	8 x 35	4.258	1.969	0.037	8.9	91.3	2.04e-09	23
25600	9 x 35	9.262	4.205	0.076	17.6	90.3	1.62e-09	23
51200	10 x 35	20.431	9.238	0.153	34.4	87.9	1.44e-09	23
102400	11 x 35	45.732	21.039	0.310	68.3	87.5	1.61e-09	23

TABLE 2. Compression to the HBSID format using Algorithm 7 of the double layer integral equation described in Section 5.1. Here $\varepsilon = 10^{-9}$ and $\ell = 35$.

For this experiment, we bench-mark the new compression algorithms by comparing their speed, accuracy, and memory requirements to the compression technique based on potential theory described in [28], run at the same precision as the randomized compression scheme. The results are presented in tables 1, 2, and 3, and summarized in Figure 9.

5.2. Operator multiplication. We next apply the proposed scheme to compute the Neumann-to-Dirichlet operator for the boundary value problem

$$\begin{cases} -\Delta u(\mathbf{x}) = 0 & \mathbf{x} \in \Omega, \\ \partial_{\mathbf{n}} u(\mathbf{x}) = g(\mathbf{x}) & \mathbf{x} \in \Gamma, \end{cases} \quad (21)$$

N	T_{compress} (sec)	t_{app} (sec)	M (MB)	M/N (reals)	E	k
400	0.053	0.001	0.4	128.3	1.51e-09	38
800	0.104	0.002	0.7	119.5	2.57e-09	37
1600	0.194	0.005	1.4	113.6	3.57e-09	35
3200	0.352	0.009	2.7	109.9	4.58e-09	34
6400	0.674	0.019	5.3	108.2	7.42e-09	33
12800	1.368	0.041	10.5	107.1	2.02e-08	31
25600	2.664	0.079	20.8	106.7	1.89e-08	29
51200	5.308	0.162	41.5	106.3	2.51e-08	28
102400	10.758	0.327	82.9	106.1	3.95e-08	25

TABLE 3. Compression to the HBSID format of the double layer integral equation described in Section 5.1, using the technique based on potential theory of [28] with $\varepsilon = 10^{-9}$.

where Γ is again the contour shown in Figure 8, where Ω is the domain exterior to Γ , and where \mathbf{n} is the unit normal vector pointing in the outwards direction from Γ . With u the solution of (21), let f denote the restriction of u to Γ , and let T denote the linear operator $T : g \mapsto f$, known as the *Neumann-to-Dirichlet (NtD) operator*. It is well-known (see Remark 6) that T can be built explicitly as the product

$$T = S \left(\frac{1}{2} I + D^* \right)^{-1}, \quad (22)$$

where S is the *single-layer operator* $[Sq](\mathbf{x}) = \int_{\Gamma} -\frac{1}{2\pi} \log |\mathbf{x} - \mathbf{y}| q(\mathbf{y}) ds(\mathbf{y})$, and where D^* is the adjoint of the double-layer operator $[D^*q](\mathbf{x}) = \int_{\Gamma} \frac{\mathbf{n}(\mathbf{x}) \cdot (\mathbf{x} - \mathbf{y})}{2\pi |\mathbf{x} - \mathbf{y}|^2} q(\mathbf{y}) ds(\mathbf{y})$. We form discrete approximations \mathbf{S} and \mathbf{D}^* to S and D^* , and compute $(\frac{1}{2}\mathbf{I} + \mathbf{D}^*)^{-1}$ using the techniques in [15], with 6th order Kapur-Rokhlin quadrature used to discretize the singular integral operator S . Then we can evaluate a discrete approximation \mathbf{T} to T via

$$\mathbf{T} = \mathbf{S} (\frac{1}{2}\mathbf{I} + \mathbf{D}^*)^{-1}.$$

For this example, we evaluated an additional error metric by testing the computed NtD operator to an exact solution u_{exact} to (21). The function u_{exact} is given as the potential from a collection of five randomly placed charges inside Γ , and then $\mathbf{f}_{\text{exact}}$ and $\mathbf{g}_{\text{exact}}$ are simply the evaluations of u_{exact} and its normal derivative on the quadrature nodes on Γ . Then the new error measure is given by

$$E_{\text{pot}} = \frac{\|\mathbf{f}_{\text{exact}} - \mathbf{T}_{\text{approx}} \mathbf{g}_{\text{exact}}\|_{\max}}{\|\mathbf{f}_{\text{exact}}\|_{\max}},$$

where $\|\cdot\|_{\max}$ is the maximum norm, and where $\mathbf{T}_{\text{approx}}$ is the compressed representation of T determined by the randomized sampling scheme proposed.

The numerical results are presented in tables 4 and 5, and summarized in Figure 10.

Remark 6. *The formula (22) for the NtD operator is derived as follows: We first look for a solution to (21) of the form $u = Sq$. Then it can be shown that q must satisfy $(1/2)q + D^*q = g$. Solving for q and using $f = Sq$, we obtain (22).*

5.3. Dimensional reduction in the Fast Multipole Method. In this section, we investigate a numerical example where \mathbf{A} is a potential evaluation map for a set of electric charges in the plane. To be precise, we let $\{\mathbf{x}_i\}_{i=1}^N$ denote a set of points located as shown in Figure 11. Then \mathbf{A} is the $N \times N$ matrix with entries

$$\mathbf{A}(i, i) = \begin{cases} -\frac{1}{2\pi} \log |\mathbf{x}_i - \mathbf{x}_j|, & \text{when } i \neq j, \\ 0 & \text{when } i = j. \end{cases}$$

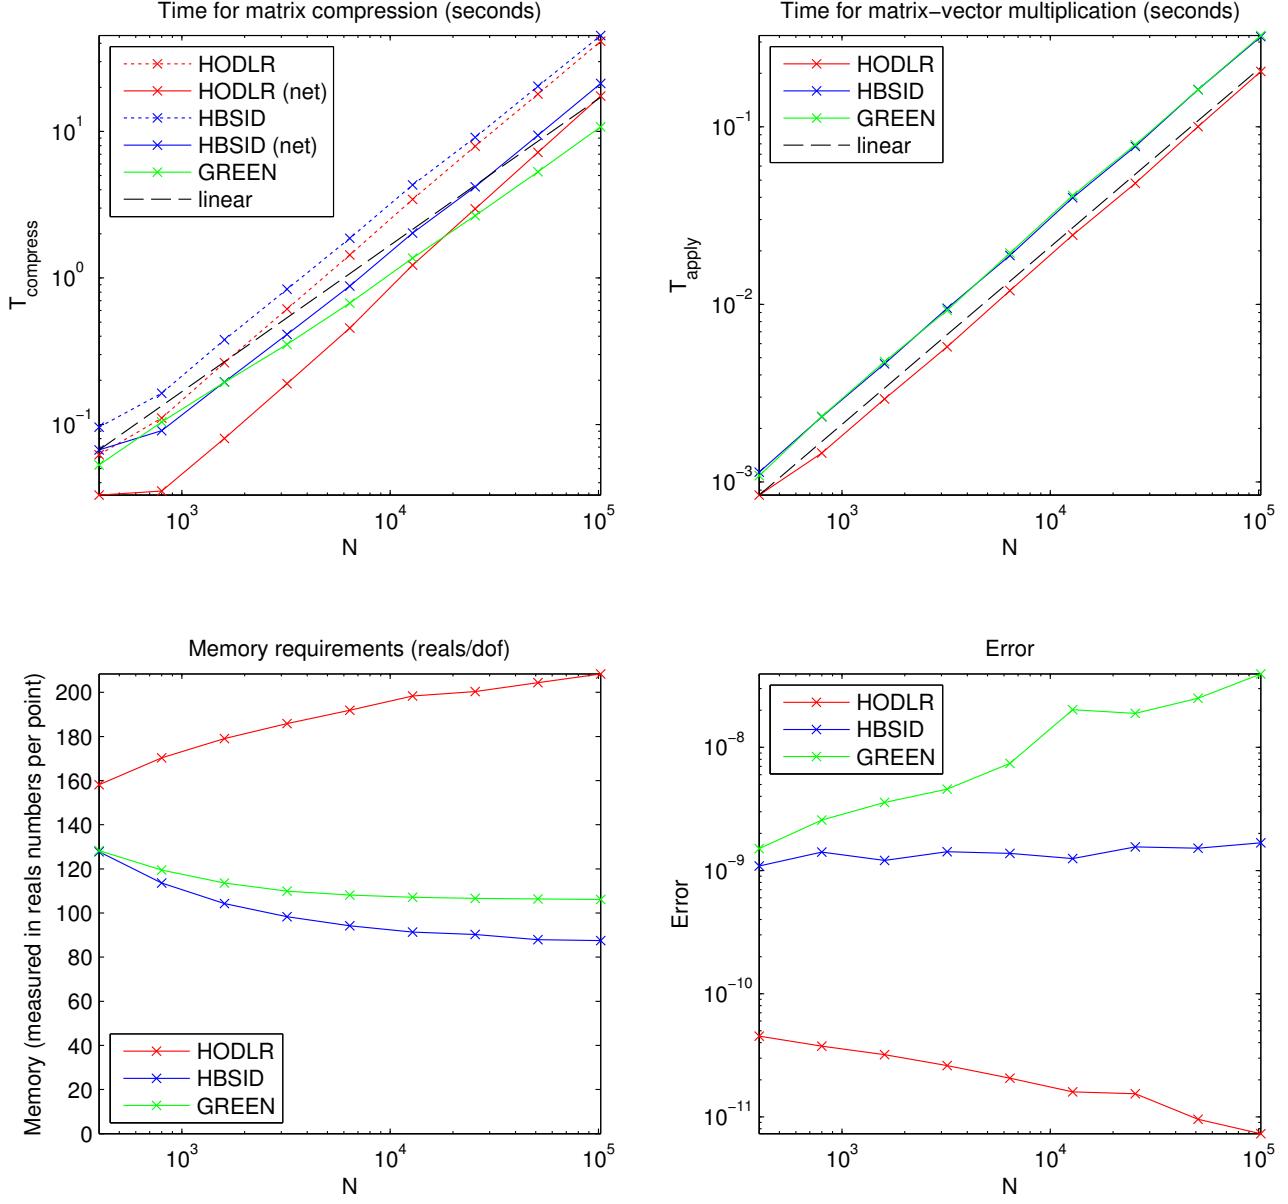


FIGURE 9. Visualization of the results presented in Tables 1, 2, and 3, pertaining to the example in Section 5.1.

We apply \mathbf{A} rapidly using the classical Fast Multipole Method [17] with 30th order expansions, which ensures that the error in the black-box code is far smaller than our requested precision of $\varepsilon = 10^{-9}$. Our FMM is implemented in Matlab as described in [27]. This implementation is quite inefficient, but has the ability to apply \mathbf{A} to an $N \times \ell$ matrix rather than single vector.

For this example, we use an additional error measure E_{pot} that compares the compressed matrix $\mathbf{A}_{\text{compress}}$ to the exact matrix \mathbf{A} , evaluated at a subset of the target points. To be precise, we picked at random a subset of 10 target locations, marked by the index vector $I \subset \{1, 2, \dots, N\}$. We then drew a sequence of ten vectors $\{\omega_i\}_{i=1}^{10}$ in which each entry is drawn at random from a uniform distribution on the interval $[0, 1]$ (observe that

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	E_{pot}	k
400	2 x 60	0.069	0.030	0.001	0.6	200.3	2.06e-11	3.87e-07	32
800	3 x 60	0.123	0.034	0.001	1.5	237.7	4.35e-11	8.22e-09	35
1600	4 x 60	0.323	0.088	0.002	3.5	283.1	4.00e-11	3.06e-09	40
3200	5 x 60	0.839	0.236	0.005	8.0	327.7	7.18e-11	9.74e-09	44
6400	6 x 60	2.093	0.657	0.011	18.1	370.0	1.08e-10	8.37e-09	45
12800	7 x 60	5.108	1.734	0.024	40.7	417.0	1.61e-10	1.76e-08	49
25600	8 x 60	12.302	4.516	0.055	89.3	457.2	3.09e-10	3.41e-08	52
51200	9 x 60	29.380	11.636	0.120	195.8	501.2	5.44e-10	7.22e-08	54
102400	10 x 60	67.879	29.072	0.284	426.0	545.3	1.11e-09	3.39e-07	55

TABLE 4. Compression to the HODLR format using Algorithm 4 of the NtD operator described in Section 5.2. Here $\varepsilon = 10^{-9}$ and $\ell = 60$.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	E_{pot}	k
400	2 x 60	0.093	0.062	0.001	0.6	211.7	3.48e-10	3.86e-07	29
800	3 x 60	0.190	0.103	0.001	1.2	198.7	2.70e-10	8.25e-09	31
1600	4 x 60	0.482	0.245	0.003	2.4	192.7	5.32e-10	5.42e-09	33
3200	5 x 60	1.154	0.557	0.006	4.6	189.9	1.57e-09	1.01e-08	36
6400	6 x 60	2.708	1.269	0.012	9.1	186.5	2.02e-09	1.35e-08	37
12800	7 x 60	6.310	2.906	0.023	18.1	185.4	3.98e-09	2.51e-08	40
25600	8 x 60	14.495	6.511	0.048	35.6	182.3	6.80e-09	4.72e-08	41
51200	9 x 60	32.906	15.033	0.094	70.8	181.4	1.47e-08	8.04e-08	42
102400	10 x 60	82.238	37.772	0.189	139.8	178.9	2.51e-08	3.46e-07	44

TABLE 5. Compression to the HODLR format using Algorithm 4 of the NtD operator described in Section 5.2. Here $\varepsilon = 10^{-9}$ and $\ell = 60$.

every charge is *positive*). Then E_{pot} is defined via

$$E_{\text{pot}} = \max_{1 \leq i \leq 10} \frac{\|\mathbf{A}(I, :) \boldsymbol{\omega}_i - \mathbf{A}_{\text{compressed}}(I, :) \boldsymbol{\omega}_i\|}{\|\mathbf{A}(I, :) \boldsymbol{\omega}_i\|}.$$

The numerical results are presented in tables 6 and 7, and summarized in Figure 12.

5.4. Compression of frontal matrices in nested dissection. Our final example applies the proposed compression schemes to the problem of constructing $O(N)$ direct solvers for the sparse linear systems arising upon the discretization of elliptic PDEs via finite difference or finite element methods. The idea is to build a solver on the classical “nested dissection” scheme of George [13, 11, 10]. In standard implementations, the problem of this direct solver is that it requires the inversion or LU-factorization of a set of successively larger dense matrices. However, it has recently been demonstrated that while these matrices are dense, they have internal structure that allows for linear or close to linear time matrix algebra to be executed, [24, 34, 30, 14]. In this manuscript, we test the proposed compression scheme on a set of matrices whose behavior is directly analogous to the matrices encountered in the algorithms of [24, 34, 30, 14]. To be precise, let \mathbf{B} denote the sparse coefficient matrix associated with a grid conduction problem on the grid shown in Figure 13. Each bar has a conductivity that is drawn at random from a uniform distribution on the interval $[1, 2]$. Let I_1, I_2, I_3 denote three index vectors that mark the three regions shown in Figure 13, set

$$\mathbf{B}_{ij} = \mathbf{B}(I_i, I_j), \quad i, j = 1, 2, 3,$$

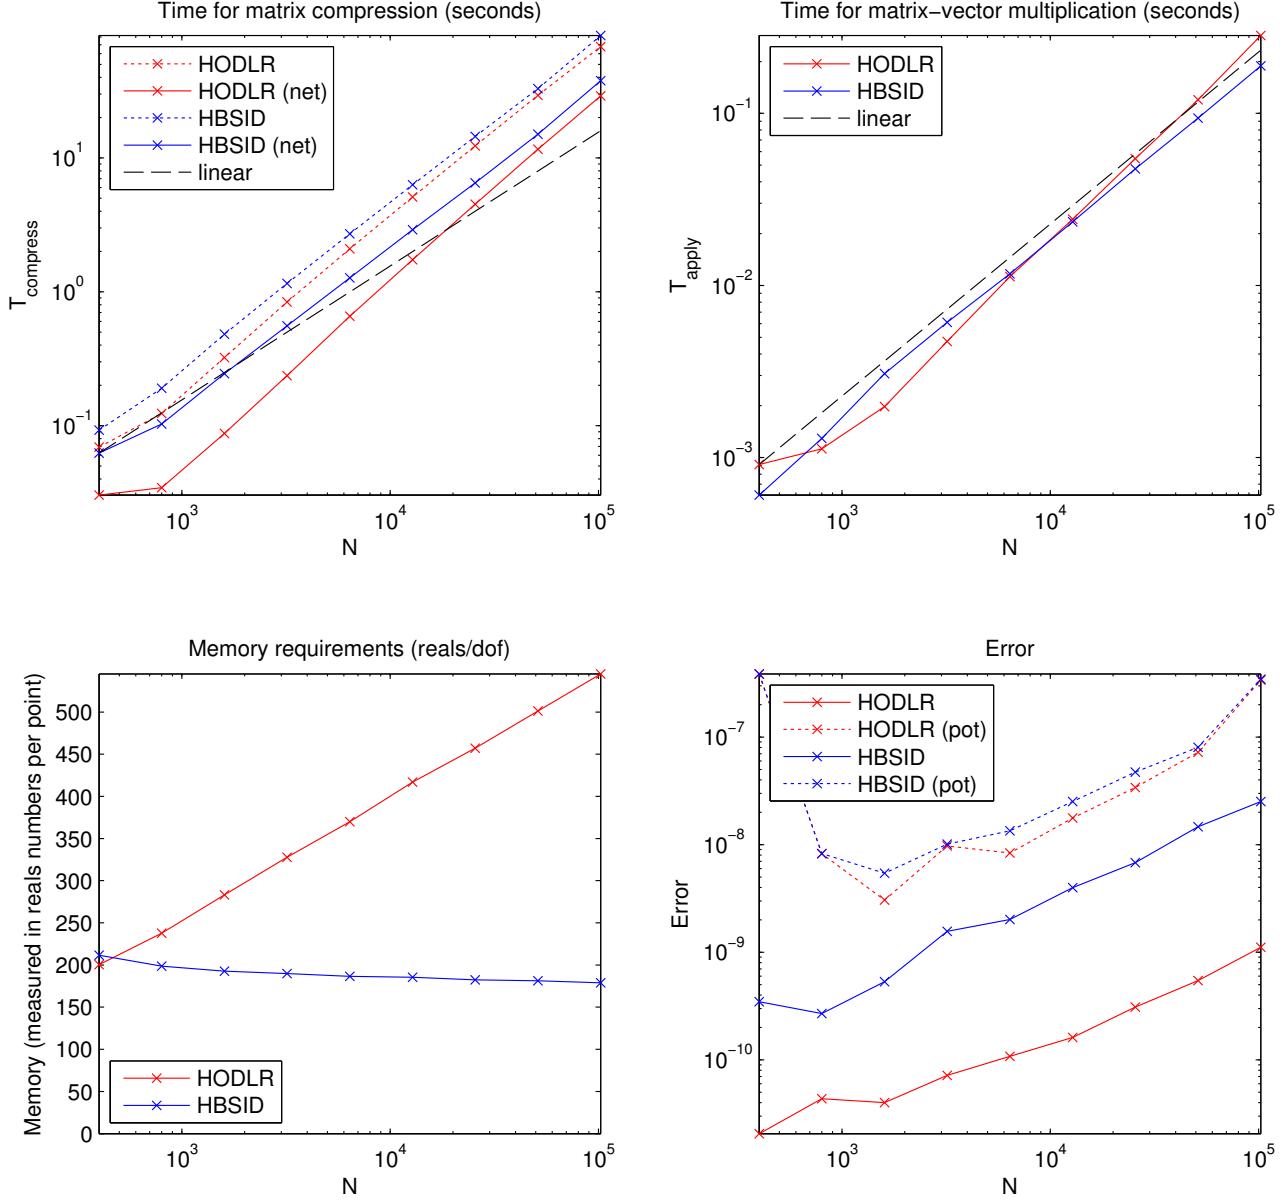


FIGURE 10. Visualization of the results presented in Tables 4, and 5, pertaining to the example in Section 5.2.

and then define the $N \times N$ matrix \mathbf{A} via

$$\mathbf{A} = \mathbf{B}_{33} - \mathbf{B}_{31}\mathbf{B}_{11}^{-1}\mathbf{B}_{13} - \mathbf{B}_{32}\mathbf{B}_{22}^{-1}\mathbf{B}_{23}. \quad (23)$$

The relevance of the matrix \mathbf{A} is discussed in some detail in Remark 7.

In our numerical experiments, the black-box application of \mathbf{A} , as defined by (23) was executed using the sparse matrix built-in routines in Matlab, which relies on UMFPACK [9] for the sparse solves implicit in the application of \mathbf{B}_{11}^{-1} and \mathbf{B}_{22}^{-1} .

The numerical results are presented in tables 8 and 9, and summarized in Figure 14.

Remark 7. *To illustrate the connection between the matrix \mathbf{A} , as defined by (23), and the LU-factorization of a matrix such as \mathbf{B} , observe first that the blocks \mathbf{B}_{12} and \mathbf{B}_{21} are zero, so that (up to a permutation of the rows*

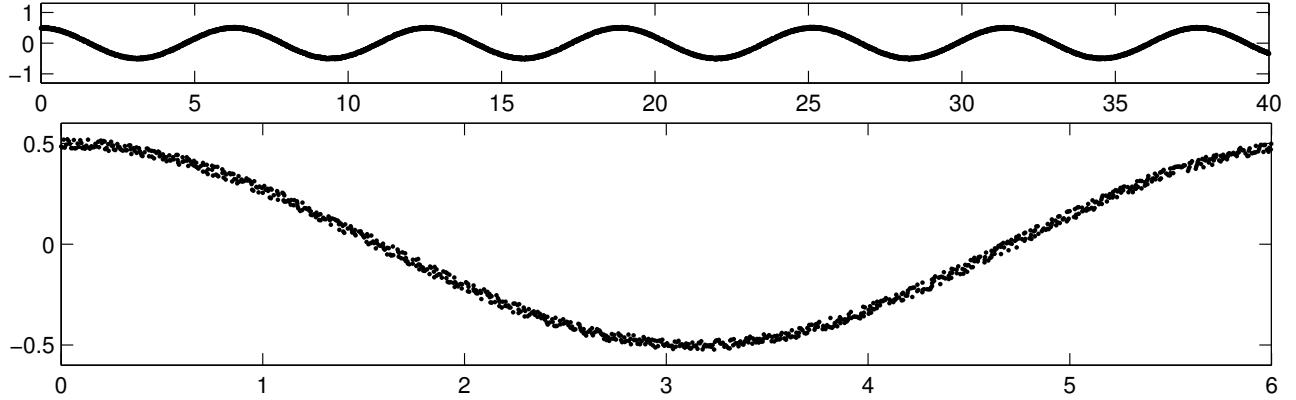


FIGURE 11. Geometry of the potential evaluation map in Section 5.3. The top figure shows the entire geometry, and the lower figure shows a magnification of a small part.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	E_{pot}	k
400	3 x 45	0.101	0.021	0.001	0.5	173.9	1.57e-10	1.16e-10	21
800	4 x 45	0.265	0.055	0.001	1.3	215.9	3.69e-10	3.52e-10	23
1600	5 x 45	0.812	0.133	0.003	3.1	257.8	1.09e-10	1.46e-10	24
3200	6 x 45	1.925	0.321	0.007	7.4	304.9	6.17e-11	8.77e-11	26
6400	7 x 45	4.509	0.783	0.015	17.2	351.5	4.22e-11	5.61e-11	28
12800	8 x 45	10.980	1.941	0.032	39.1	400.6	3.11e-11	3.63e-11	30
25600	9 x 45	26.271	4.788	0.068	87.1	446.0	2.56e-11	3.38e-11	30
51200	10 x 45	62.119	11.938	0.149	192.5	492.7	2.26e-11	2.67e-11	32
102400	11 x 45	158.617	30.621	0.362	419.5	537.0	1.96e-11	2.19e-11	32

TABLE 6. Compression to the HODLR format using Algorithm 4 of the potential evaluation matrix described in Section 5.3. Here $\varepsilon = 10^{-9}$ and $\ell = 45$.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	E_{pot}	k
400	3 x 45	0.140	0.062	0.001	0.5	164.1	3.84e-09	2.96e-09	31
800	4 x 45	0.358	0.149	0.002	1.0	164.8	2.29e-08	2.10e-08	33
1600	5 x 45	1.007	0.329	0.005	2.0	166.1	1.12e-08	1.28e-08	37
3200	6 x 45	2.306	0.713	0.010	4.1	166.5	1.29e-08	1.08e-08	39
6400	7 x 45	5.334	1.554	0.020	8.0	164.7	8.47e-09	5.13e-09	40
12800	8 x 45	12.366	3.350	0.040	15.8	161.7	7.75e-09	4.66e-09	42
25600	9 x 45	28.813	7.170	0.081	30.9	158.1	7.28e-09	7.56e-09	43
51200	10 x 45	64.838	15.572	0.166	60.1	154.0	1.47e-08	9.91e-09	44
102400	11 x 45	164.750	36.686	0.335	116.7	149.4	1.23e-08	1.03e-08	43

TABLE 7. Compression to the HBSID format using Algorithm 7 of the potential evaluation matrix described in Section 5.3. Here $\varepsilon = 10^{-9}$ and $\ell = 45$.

and columns),

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{0} & \mathbf{B}_{13} \\ \mathbf{0} & \mathbf{B}_{22} & \mathbf{B}_{23} \\ \mathbf{B}_{31} & \mathbf{B}_{32} & \mathbf{B}_{33} \end{bmatrix}.$$

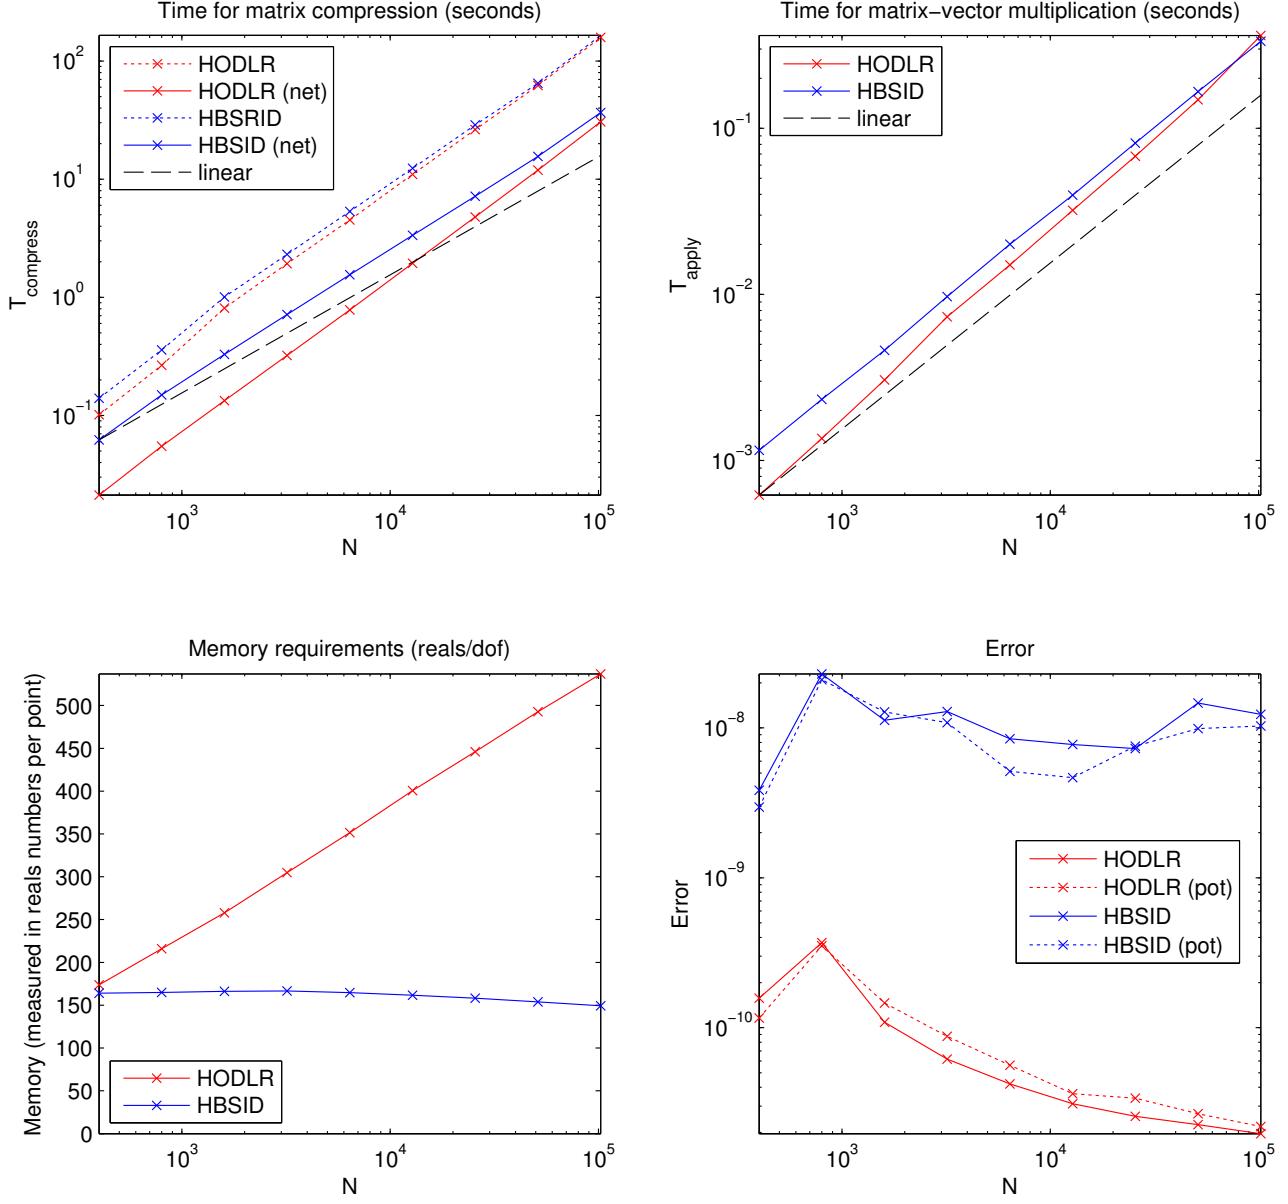


FIGURE 12. Visualization of the results presented in Tables 6, and 7, pertaining to the example in Section 5.3.

Next suppose that we can somehow determine the LU-factorizations of \mathbf{B}_{11} and \mathbf{B}_{22} ,

$$\mathbf{B}_{11} = \mathbf{L}_{11} \mathbf{U}_{11}, \quad \text{and} \quad \mathbf{B}_{22} = \mathbf{L}_{22} \mathbf{U}_{22}.$$

Then the LU-factorization of \mathbf{B} is given by

$$\mathbf{B} = \begin{bmatrix} \mathbf{L}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_{22} & \mathbf{0} \\ \mathbf{B}_{31} \mathbf{U}_{11}^{-1} & \mathbf{B}_{32} \mathbf{U}_{22}^{-1} & \mathbf{L}_{33} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{11} & \mathbf{0} & \mathbf{L}_{11}^{-1} \mathbf{B}_{13} \\ \mathbf{0} & \mathbf{U}_{22} & \mathbf{L}_{22}^{-1} \mathbf{B}_{23} \\ \mathbf{0} & \mathbf{0} & \mathbf{U}_{33} \end{bmatrix},$$

where \mathbf{L}_{33} and \mathbf{U}_{33} are defined as the LU factors of

$$\mathbf{L}_{33} \mathbf{U}_{33} = \underbrace{\mathbf{B}_{33} - \mathbf{B}_{31} \mathbf{U}_{11}^{-1} \mathbf{L}_{11}^{-1} \mathbf{B}_{13} - \mathbf{B}_{32} \mathbf{U}_{22}^{-1} \mathbf{L}_{22}^{-1} \mathbf{B}_{23}}_{=: \mathbf{A}}. \quad (24)$$

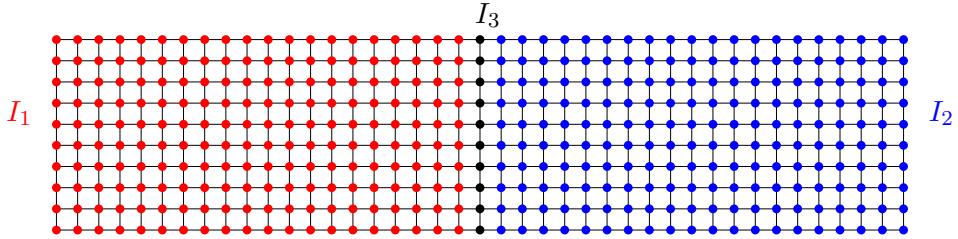


FIGURE 13. Geometry of problem described in Section 5.4. We consider a grid conduction problem on the grid shown. As N is increased, the *width* or the grid is fixed at 41 nodes, while the *height* of the grid equals N .

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	k
400	3 x 25	0.348	0.012	0.001	0.4	115.0	1.83e-14	9
800	4 x 25	0.849	0.025	0.001	0.8	133.4	1.47e-14	9
1600	5 x 25	2.143	0.064	0.002	1.9	151.6	1.40e-14	9
3200	6 x 25	5.588	0.141	0.005	4.1	169.7	1.43e-14	9
6400	7 x 25	13.926	0.336	0.011	9.2	187.8	1.37e-14	9
12800	8 x 25	38.558	0.903	0.023	20.1	205.8	1.36e-14	9
25600	9 x 25	91.382	2.322	0.048	43.7	223.8	1.29e-14	9
51200	10 x 25	208.704	5.874	0.098	94.5	241.8	1.32e-14	9
102400	11 x 25	468.981	14.203	0.204	203.0	259.8	1.30e-14	9

TABLE 8. Compression to the HODLR format using Algorithm 4 of a simulated “frontal matrix” in the nested dissection technique, as described in Section 5.4. Here $\varepsilon = 10^{-9}$ and $\ell = 25$.

N	N_{matvec}	T_{compress} (sec)	T_{net} (sec)	T_{app} (sec)	M (MB)	M/N (reals)	E	k
400	3 x 25	0.361	0.032	0.001	0.4	121.0	3.80e-14	18
800	4 x 25	0.867	0.066	0.002	0.8	124.9	4.75e-14	18
1600	5 x 25	2.180	0.150	0.005	1.6	127.7	4.34e-14	18
3200	6 x 25	5.575	0.322	0.010	3.2	129.5	4.27e-14	18
6400	7 x 25	14.470	0.708	0.019	6.4	130.6	4.30e-14	18
12800	8 x 25	38.945	1.525	0.040	12.8	131.3	4.19e-14	18
25600	9 x 25	91.563	3.253	0.081	25.7	131.7	4.11e-14	18
51200	10 x 25	208.145	7.440	0.157	51.5	131.9	4.10e-14	18
102400	11 x 25	465.309	15.759	0.314	103.1	132.0	4.07e-14	18

TABLE 9. Compression to the HBSID format using Algorithm 7 of a simulated “frontal matrix” in the nested dissection technique, as described in Section 5.4. Here $\varepsilon = 10^{-9}$ and $\ell = 25$.

Observe that since the matrices \mathbf{L}_{11} , \mathbf{U}_{11} , \mathbf{L}_{22} , \mathbf{U}_{22} are all triangular, their inverses are inexpensive to apply. To summarize, if one can cheaply evaluate the LU factorization (24), then the task of LU factoring \mathbf{B} directly reduces to the task of LU factoring the two matrices \mathbf{B}_{11} and \mathbf{B}_{22} , which both involve about half as many variables as \mathbf{B} . The classical nested dissection idea is now to apply this observation recursively to factor \mathbf{B}_{11} and \mathbf{B}_{22} . The problem of this scheme has traditionally been that in order to evaluate \mathbf{L}_{33} and \mathbf{U}_{33} in (24), one must factorize the dense matrix \mathbf{A} .

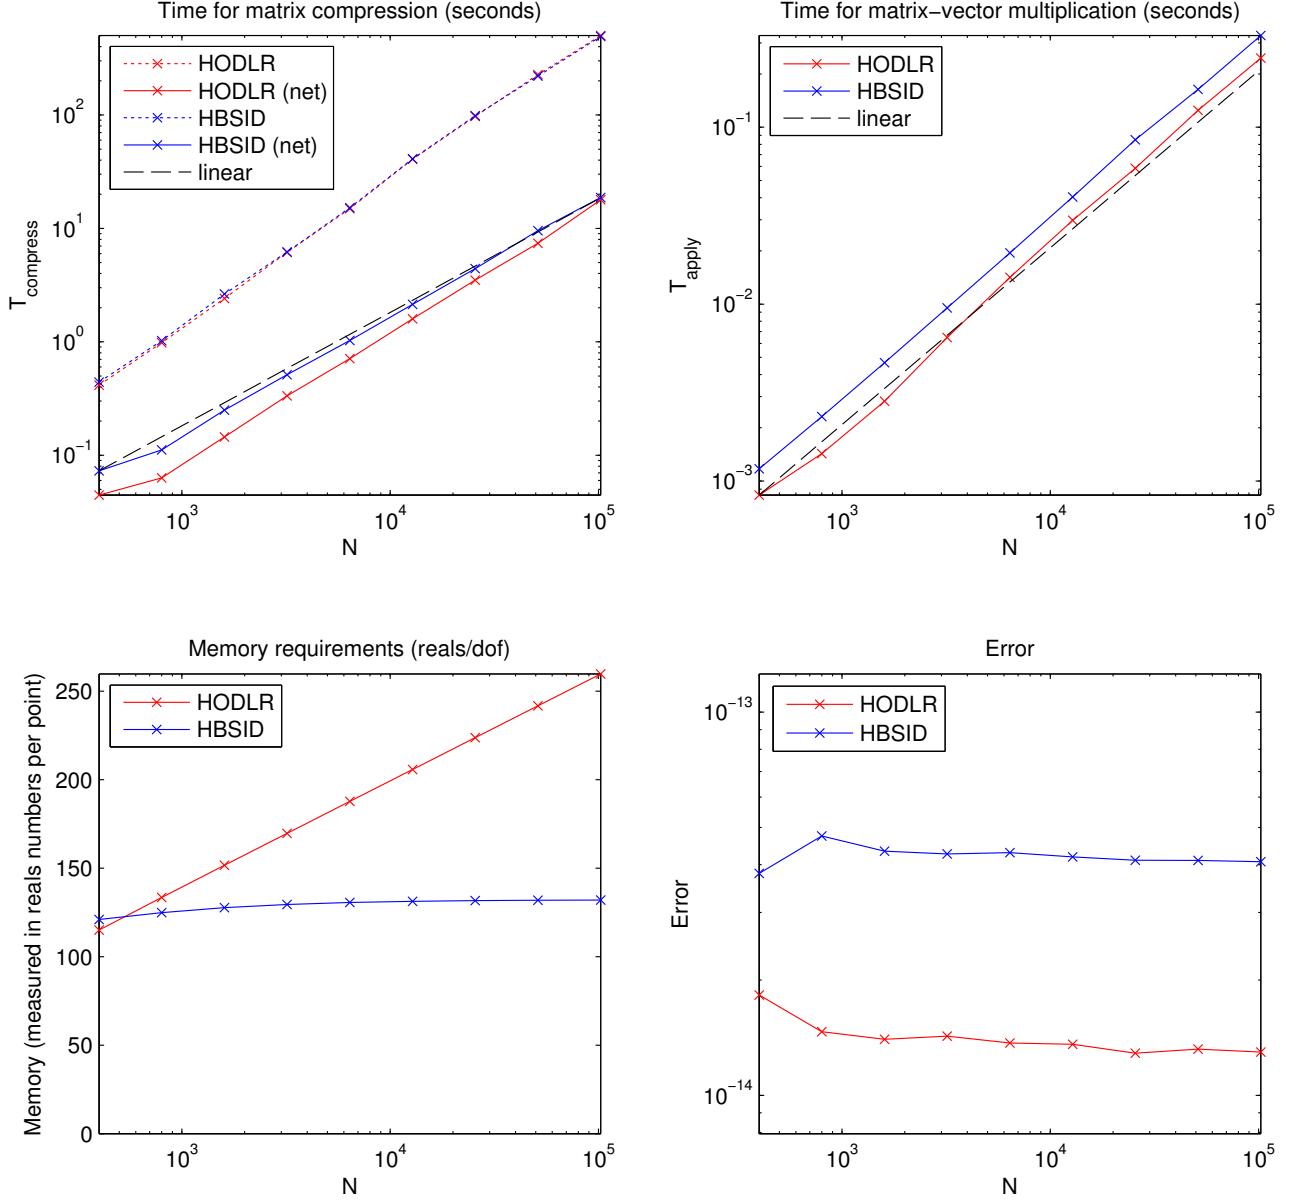


FIGURE 14. Visualization of the results presented in Tables 8, and 9, pertaining to the example in Section 5.4.

5.5. Summary of observations from numerical experiments. To close this section, we make some observations and conjectures:

- In all examples examined, numerical evidence supports the claims on asymptotic scaling made in Sections 3.1 and 4.4.
- Excellent approximation errors are obtained in every case. Aggregation of errors over levels is a very minor problem.
- The computational time is in all cases dominated by the time spent in the external black-box multiplier. As a consequence, the primary route by which the proposed algorithm could be improved would be to lessen the number of matrix-vector multiplications required.

Build outgoing expansions on level m .

loop over all nodes τ on level m

$$\hat{\mathbf{q}}_\tau = \mathbf{V}_\tau^* \mathbf{q}(I_\tau)$$

end loop

Build outgoing expansions on levels coarser than m (upwards pass).

loop over levels $\ell = (m-1) : (-1) : 1$

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of τ .

$$\hat{\mathbf{q}}_\tau = \mathbf{V}_\tau^* \begin{bmatrix} \hat{\mathbf{q}}_\alpha \\ \hat{\mathbf{q}}_\beta \end{bmatrix}.$$

end loop

end loop

Build incoming expansions for the children of the root.

Let $\{\alpha, \beta\}$ denote the children of the root node.

$$\hat{\mathbf{u}}_\alpha = \mathbf{B}_{\alpha, \beta} \hat{\mathbf{q}}_\beta$$

$$\hat{\mathbf{u}}_\beta = \mathbf{B}_{\beta, \alpha} \hat{\mathbf{q}}_\alpha.$$

Build incoming expansions on levels coarser than m (downwards pass).

loop over levels $\ell = (m-1) : (-1) : 1$

loop over boxes τ on level ℓ

Let $\{\alpha, \beta\}$ denote the children of τ .

$$\hat{\mathbf{u}}_\alpha = \mathbf{B}_{\alpha, \beta} \hat{\mathbf{q}}_\beta + \mathbf{U}_\tau(J_\alpha, :) \hat{\mathbf{u}}_\tau.$$

$$\hat{\mathbf{u}}_\beta = \mathbf{B}_{\beta, \alpha} \hat{\mathbf{q}}_\alpha + \mathbf{U}_\tau(J_\beta, :) \hat{\mathbf{u}}_\tau.$$

end loop

end loop

Build incoming expansions on level m .

loop over boxes τ on level m

$$\mathbf{u}(I_\tau) = \mathbf{U}_\tau \hat{\mathbf{u}}_\tau$$

end loop

FIGURE 15. Application of $\mathbf{A}^{(m)}$ in the HBS framework. Given a vector \mathbf{q} of charges, build the vector $\mathbf{u} = \mathbf{A}^{(m)} \mathbf{q}$ of potentials.

- For modest problem sizes, the HODLR algorithm is very fast and easy to use. However, as problems grow large, the memory requirements of the HODLR format become slightly problematic, and the HBSID algorithm gains a more pronounced advantage.

REFERENCES

- [1] S. AMBIKASARAN AND E. DARVE, *An $o(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices*, Journal of Scientific Computing, 57 (2013), pp. 477–501.
- [2] M. BEBENDORF, *Approximation of boundary element matrices*, Numerische Mathematik, 86 (2000), pp. 565–589.
- [3] ———, *Hierarchical matrices*, vol. 63 of Lecture Notes in Computational Science and Engineering, Springer-Verlag, Berlin, 2008. A means to efficiently solve elliptic boundary value problems.
- [4] S. BÖRM, *Efficient numerical methods for non-local operators*, vol. 14 of EMS Tracts in Mathematics, European Mathematical Society (EMS), Zürich, 2010. \mathcal{H}^2 -matrix compression, algorithms and analysis.
- [5] S. BÖRM AND L. GRASEDYCK, *Hybrid cross approximation of integral operators*, Numerische Mathematik, 101 (2005), pp. 221–249.
- [6] T. F. CHAN, *Rank revealing qr factorizations*, Linear Algebra and its Applications, 88-89 (1987), pp. 67 – 82.

- [7] S. CHANDRASEKARAN, M. GU, AND W. LYONS, *A fast adaptive solver for hierarchically semiseparable representations*, Calcolo, 42 (2005), pp. 171–185.
- [8] H. CHENG, Z. GIMBUTAS, P. MARTINSSON, AND V. ROKHLIN, *On the compression of low rank matrices*, SIAM Journal of Scientific Computing, 26 (2005), pp. 1389–1404.
- [9] T. A. DAVIS, *Algorithm 832: Umfpack v4. 3—an unsymmetric-pattern multifrontal method*, ACM Transactions on Mathematical Software (TOMS), 30 (2004), pp. 196–199.
- [10] ———, *Direct methods for sparse linear systems*, vol. 2, Siam, 2006.
- [11] I. DUFF, A. ERISMAN, AND J. REID, *Direct Methods for Sparse Matrices*, Oxford, 1989.
- [12] K. FREDERIX AND M. V. BAREL, *Solving a large dense linear system by adaptive cross approximation*, Journal of Computational and Applied Mathematics, 234 (2010), pp. 3181 – 3195.
- [13] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. on Numerical Analysis, 10 (1973), pp. 345–363.
- [14] A. GILLMAN AND P. MARTINSSON, *A direct solver with $o(n)$ complexity for variable coefficient elliptic pdes discretized via a high-order composite spectral collocation method*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2023–A2046.
- [15] A. GILLMAN, P. YOUNG, AND P.-G. MARTINSSON, *A direct solver $o(n)$ complexity for integral equations on one-dimensional domains*, Frontiers of Mathematics in China, 7 (2012), pp. 217–247. 10.1007/s11464-012-0188-3.
- [16] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, third ed., 1996.
- [17] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.
- [18] M. GU AND S. C. EISENSTAT, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869.
- [19] W. HACKBUSCH, *A sparse matrix arithmetic based on H-matrices; Part I: Introduction to H-matrices*, Computing, 62 (1999), pp. 89–108.
- [20] W. HACKBUSCH AND S. BÖRM, *Data-sparse approximation by adaptive \mathcal{H}^2 -matrices*, Computing, 69 (2002), pp. 1–35.
- [21] W. HACKBUSCH, B. KHOROMSKIJ, AND S. SAUTER, *On \mathcal{H}^2 -matrices*, in *Lectures on Applied Mathematics*, Springer Berlin, 2002, pp. 9–29.
- [22] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.
- [23] D. HUYBRECHS, *Multiscale and hybrid methods for the solution of oscillatory integral equations*, PhD thesis, Katholieke Universiteit Leuven, 2006.
- [24] S. LE BORNE, L. GRASEDYCK, AND R. KRIEMANN, *Domain-decomposition based \mathcal{H} -LU preconditioners*, in *Domain decomposition methods in science and engineering XVI*, vol. 55 of *Lect. Notes Comput. Sci. Eng.*, Springer, Berlin, 2007, pp. 667–674.
- [25] L. LIN, J. LU, AND L. YING, *Fast construction of hierarchical matrix representation from matrixvector multiplication*, Journal of Computational Physics, 230 (2011), pp. 4071 – 4087.
- [26] P. MARTINSSON, *Rapid factorization of structured matrices via randomized sampling*, 2008. arXiv:0806.2339.
- [27] ———, *A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix*, SIAM Journal on Matrix Analysis and Applications, 32 (2011), pp. 1251–1274.
- [28] P. MARTINSSON AND V. ROKHLIN, *A fast direct solver for boundary integral equations in two dimensions*, J. Comp. Phys., 205 (2005), pp. 1–23.
- [29] ———, *An accelerated kernel independent fast multipole method in one dimension*, SIAM Journal of Scientific Computing, 29 (2007), pp. 1160–11178.
- [30] P.-G. MARTINSSON, *A fast direct solver for a class of elliptic partial differential equations*, J. Sci. Comput., 38 (2009), pp. 316–330.
- [31] P.-G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *A randomized algorithm for the approximation of matrices*, Tech. Report Yale CS research report YALEU/DCS/RR-1361, Yale University, Computer Science Department, 2006.
- [32] ———, *A randomized algorithm for the decomposition of matrices*, Appl. Comput. Harmon. Anal., 30 (2011), pp. 47–68.
- [33] J. XIA, *Randomized sparse direct solvers*, SIAM Journal on Matrix Analysis and Applications, 34 (2013), pp. 197–227.
- [34] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Superfast multifrontal method for large structured linear systems of equations*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1382–1411.
- [35] ———, *Fast algorithms for hierarchically semiseparable matrices*, Numerical Linear Algebra with Applications, 17 (2010), pp. 953–976.